Cisco ASA Series ASDM Configuration Guide

Software Version 7.0 for the ASA 5505, ASA 5510, ASA 5520, ASA 5540, ASA 5550, ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X, ASA 5580, ASA 5585-X, and the ASA Services Module

Released: October 29, 2012
Updated: February 25, 2013

Cisco Systems, Inc.
www.cisco.com

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

Text Part Number: N/A, Online only

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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

Cisco ASA Series ASDM Configuration Guide
Copyright © 2012-2013 Cisco Systems, Inc. All rights reserved.

About This Guide lxi

Document Objectives lxi

Audience lxi

Related Documentation lxii

Conventions lxii

Obtaining Documentation and Submitting a Service Request lxiii

PART 1 **Getting Started with the ASA**

CHAPTER 1 **Introduction to the Cisco ASA** 1-1

ASDM Client Operating System and Browser Requirements 1-1

Hardware and Software Compatibility 1-4

VPN Specifications 1-4

New Features 1-5

New Features in ASA 8.4(5)/ASDM 7.0(2) 1-5

New Features in ASA 9.0(1)/ASDM 7.0(1) 1-5

How the ASA Services Module Works with the Switch 1-19

Firewall Functional Overview 1-20

Security Policy Overview 1-21

Firewall Mode Overview 1-23

Stateful Inspection Overview 1-24

VPN Functional Overview 1-25

Security Context Overview 1-25

ASA Clustering Overview 1-26

CHAPTER 2 **Configuring the Switch for Use with the ASA Services Module** 2-1

Information About the Switch 2-1

Guidelines and Limitations 2-2

Verifying the Module Installation 2-3

Assigning VLANs to the ASA Services Module 2-4

Using the MSFC as a Directly Connected Router 2-5

Information About SVIs 2-6

Configuring SVIs 2-8

Configuring the Switch for ASA Failover 2-9
Assigning VLANs to the Secondary ASA Services Module 2-10
Adding a Trunk Between a Primary Switch and Secondary Switch 2-10
Ensuring Compatibility with Transparent Firewall Mode 2-10
Enabling Autostate Messaging for Rapid Link Failure Detection 2-10
Resetting the ASA Services Module 2-11
Monitoring the ASA Services Module 2-11
Feature History for the Switch for Use with the ASA Services Module 2-13

CHAPTER 3 **Getting Started** 3-1
Accessing the Appliance Command-Line Interface 3-1
Accessing the ASA Services Module Command-Line Interface 3-2
Logging Into the ASA Services Module 3-2
Logging Out of a Console Session 3-4
Logging Out of a Telnet Session 3-6
Configuring ASDM Access for Appliances 3-6
Accessing ASDM Using the Factory Default Configuration 3-6
Accessing ASDM Using a Non-Default Configuration (ASA 5505) 3-7
Accessing ASDM Using a Non-Default Configuration (ASA 5510 and Higher) 3-9
Configuring ASDM Access for the ASA Services Module 3-11
Starting ASDM 3-13
Connecting to ASDM for the First Time 3-14
Starting ASDM from the ASDM-IDM Launcher 3-15
Starting ASDM from the Java Web Start Application 3-15
Using ASDM in Demo Mode 3-16
Factory Default Configurations 3-17
Restoring the Factory Default Configuration 3-18
ASA 5505 Default Configuration 3-20
ASA 5510 and Higher Default Configuration 3-24
Getting Started with the Configuration 3-24
Using the Command Line Interface Tool in ASDM 3-25
Using the Command Line Interface Tool 3-25
Handling Command Errors 3-26
Using Interactive Commands 3-26
Avoiding Conflicts with Other Administrators 3-26
Showing Commands Ignored by ASDM on the Device 3-26
Applying Configuration Changes to Connections 3-27

CHAPTER 4 **Using ASDM** 4-1

Information About the ASDM User Interface 4-1

Navigating in the ASDM User Interface 4-3

Menus 4-4

File Menu 4-4

View Menu 4-5

Tools Menu 4-6

Wizards Menu 4-8

Window Menu 4-9

Help Menu 4-9

Toolbar 4-10

ASDM Assistant 4-10

Status Bar 4-11

Connection to Device 4-11

Device List 4-11

Common Buttons 4-12

Keyboard Shortcuts 4-12

Find Function 4-14

Using the Find Function in Most ASDM Panes 4-14

Using the Find Function in the ACL Manager Pane 4-15

Enabling Extended Screen Reader Support 4-16

Organizational Folder 4-16

About the Help Window 4-16

Header Buttons 4-17

Browser Window 4-17

Home Pane (Single Mode and Context) 4-17

Device Dashboard Tab 4-18

Firewall Dashboard Tab 4-22

Cluster Dashboard Tab 4-25

Cluster Firewall Dashboard Tab 4-26

Content Security Tab 4-27

Intrusion Prevention Tab 4-28

ASA CX Status Tab 4-30

Home Pane (System) 4-31

Defining ASDM Preferences 4-32

Using the ASDM Assistant 4-33

Enabling History Metrics 4-34

Unsupported Commands 4-34
Ignored and View-Only Commands 4-35
Effects of Unsupported Commands 4-35
Discontinuous Subnet Masks Not Supported 4-36
Interactive User Commands Not Supported by the ASDM CLI Tool 4-36

CHAPTER 5 **Managing Feature Licenses** 5-1

Supported Feature Licenses Per Model 5-1
Licenses Per Model 5-1
License Notes 5-18
VPN License and Feature Compatibility 5-23
Information About Feature Licenses 5-23
Preinstalled License 5-24
Permanent License 5-24
Time-Based Licenses 5-24
Shared AnyConnect Premium Licenses 5-27
Failover or ASA Cluster Licenses 5-31
No Payload Encryption Models 5-33
Licenses FAQ 5-33
Guidelines and Limitations 5-34
Configuring Licenses 5-36
Obtaining an Activation Key 5-36
Activating or Deactivating Keys 5-37
Configuring a Shared License 5-38
Monitoring Licenses 5-40
Viewing Your Current License 5-40
Monitoring the Shared License 5-41
Feature History for Licensing 5-41

CHAPTER 6 **Configuring the Transparent or Routed Firewall** 6-1

Configuring the Firewall Mode 6-1
Information About the Firewall Mode 6-1
Licensing Requirements for the Firewall Mode 6-6
Default Settings 6-6
Guidelines and Limitations 6-6
Setting the Firewall Mode (Single Mode) 6-8
Feature History for Firewall Mode 6-9
Configuring ARP Inspection for the Transparent Firewall 6-10
Information About ARP Inspection 6-11

Licensing Requirements for ARP Inspection 6-11
Default Settings 6-11
Guidelines and Limitations 6-11
Configuring ARP Inspection 6-12
Feature History for ARP Inspection 6-14
Customizing the MAC Address Table for the Transparent Firewall 6-14
Information About the MAC Address Table 6-15
Licensing Requirements for the MAC Address Table 6-15
Default Settings 6-15
Guidelines and Limitations 6-15
Configuring the MAC Address Table 6-16
Feature History for the MAC Address Table 6-17
Firewall Mode Examples 6-17
How Data Moves Through the ASA in Routed Firewall Mode 6-17
How Data Moves Through the Transparent Firewall 6-23

CHAPTER 7 **Using the Startup Wizard** 7-1
Accessing the Startup Wizard 7-1
Licensing Requirements for the Startup Wizard 7-1
Guidelines and Limitations 7-1
Startup Wizard Screens 7-2
Starting Point or Welcome 7-2
Basic Configuration 7-3
Interface Screens 7-3
Static Routes 7-4
Easy VPN Remote Configuration (ASA 5505, Single Mode, Routed Mode) 7-4
DHCP Server 7-4
Address Translation (NAT/PAT) 7-5
Administrative Access 7-5
IPS Basic Configuration 7-5
Time Zone and Clock Configuration 7-5
Auto Update Server (Single Mode) 7-5
Startup Wizard Summary 7-6
Feature History for the Startup Wizard 7-6

PART 1 **Configuring High Availability and Scalability**

CHAPTER 9 **Configuring Multiple Context Mode** 9-1

Information About Security Contexts 9-1
Common Uses for Security Contexts 9-2
Context Configuration Files 9-2
How the ASA Classifies Packets 9-3
Cascading Security Contexts 9-6
Management Access to Security Contexts 9-7
Information About Resource Management 9-8
Information About MAC Addresses 9-11
Licensing Requirements for Multiple Context Mode 9-13
Prerequisites 9-14
Guidelines and Limitations 9-14
Default Settings 9-15
Configuring Multiple Contexts 9-15
Task Flow for Configuring Multiple Context Mode 9-15
Enabling or Disabling Multiple Context Mode 9-16
Configuring a Class for Resource Management 9-17
Configuring a Security Context 9-20
Automatically Assigning MAC Addresses to Context Interfaces 9-24
Changing Between Contexts and the System Execution Space 9-25
Managing Security Contexts 9-26
Removing a Security Context 9-26
Changing the Admin Context 9-27
Changing the Security Context URL 9-28
Reloading a Security Context 9-29
Monitoring Security Contexts 9-30
Monitoring Context Resource Usage 9-31
Viewing Assigned MAC Addresses 9-32
Feature History for Multiple Context Mode 9-33

CHAPTER 10 **Configuring a Cluster of ASAs** 10-1

Information About ASA Clustering 10-1
How the ASA Cluster Fits into Your Network 10-2
Performance Scaling Factor 10-2
Cluster Members 10-2
ASA Cluster Interfaces 10-4
Cluster Control Link 10-6
High Availability within the ASA Cluster 10-7

Configuration Replication 10-8
ASA Cluster Management 10-9
Load Balancing Methods 10-10
How the ASA Cluster Manages Connections 10-13
ASA Features and Clustering 10-14
Licensing Requirements for ASA Clustering 10-21
Prerequisites for ASA Clustering 10-21
Guidelines and Limitations 10-23
Default Settings 10-24
Configuring ASA Clustering 10-24
Task Flow for ASA Cluster Configuration 10-25
Cabling the Cluster Units and Configuring Upstream and Downstream Equipment 10-25
Configuring the Cluster Interface Mode on the Master Unit 10-27
Configuring the Management Interface on Each Unit 10-29
Enabling the Cluster Control Link Interface on the Master Unit 10-29
Backing Up Your Configurations (Recommended) 10-32
Configuring Interfaces on the Master Unit 10-32
Adding or Joining an ASA Cluster 10-44
Managing ASA Cluster Members 10-50
Configuring ASA Cluster Parameters 10-51
Adding a New Slave from the Master Unit 10-54
Becoming an Inactive Member 10-56
Inactivating a Slave Member from the Master Unit 10-57
Leaving the Cluster 10-58
Changing the Master Unit 10-61
Executing a Command Cluster-Wide 10-62
Monitoring the ASA Cluster 10-63
Cluster Dashboards 10-63
Monitoring Screens 10-64
Related Features 10-65
Configuration Examples for ASA Clustering 10-66
Firewall on a Stick 10-66
Traffic Segregation 10-67
Redundant Interface (PBR or ECMP) 10-67
Spanned EtherChannel With Backup Links 10-68
Feature History for ASA Clustering 10-71

CHAPTER 11 **Information About Failover** 11-1

Introduction to Failover and High Availability 11-1

Failover System Requirements 11-2

Hardware Requirements 11-2

Software Requirements 11-2

License Requirements 11-2

Failover and Stateful Failover Links 11-3

Failover Link 11-3

Stateful Failover Link 11-4

Avoiding Interrupted Failover Links 11-5

Active/Active and Active/Standby Failover 11-8

Determining Which Type of Failover to Use 11-8

Stateless (Regular) and Stateful Failover 11-9

Stateless (Regular) Failover 11-9

Stateful Failover 11-10

Intra- and Inter-Chassis Module Placement for the ASA Services Module 11-11

Intra-Chassis Failover 11-11

Inter-Chassis Failover 11-12

Transparent Firewall Mode Requirements 11-15

Auto Update Server Support in Failover Configurations 11-16

Auto Update Process Overview 11-16

Monitoring the Auto Update Process 11-17

Failover Health Monitoring 11-18

Unit Health Monitoring 11-19

Interface Monitoring 11-19

Failover Times 11-20

Failover Messages 11-20

Failover System Messages 11-20

Debug Messages 11-21

SNMP 11-21

CHAPTER 12 **Configuring Active/Standby Failover** 12-1

Information About Active/Standby Failover 12-1

Active/Standby Failover Overview 12-1

Primary/Secondary Status and Active/Standby Status 12-2

Device Initialization and Configuration Synchronization 12-2

Command Replication 12-3

Failover Triggers 12-4

Failover Actions 12-4
Optional Active/Standby Failover Settings 12-5
Licensing Requirements for Active/Standby Failover 12-5
Prerequisites for Active/Standby Failover 12-6
Guidelines and Limitations 12-6
Configuring Active/Standby Failover 12-7
Using the High Availability and Scalability Wizard 12-7
Configuring Failover (Without the Wizard) 12-10
Configuring Interface Standby Addresses 12-12
Configuring Optional Active/Standby Failover Settings 12-13
Controlling Failover 12-16
Forcing Failover 12-16
Disabling Failover 12-17
Restoring a Failed Unit 12-17
Monitoring Active/Standby Failover 12-17
Feature History for Active/Standby Failover 12-17

CHAPTER 13 **Configuring Active/Active Failover** 13-1
Information About Active/Active Failover 13-1
Active/Active Failover Overview 13-1
Primary/Secondary Status and Active/Standby Status 13-2
Device Initialization and Configuration Synchronization 13-3
Command Replication 13-3
Failover Triggers 13-4
Failover Actions 13-4
Optional Active/Active Failover Settings 13-6
Licensing Requirements for Active/Active Failover 13-6
Prerequisites for Active/Active Failover 13-6
Guidelines and Limitations 13-7
Configuring Active/Active Failover 13-8
Task Flow for Using the High Availability and Scalability Wizard 13-8
Configuring Failover (Without the Wizard) 13-12
Configuring Asymmetric Routing Groups in Multiple Context Mode 13-20
Controlling Failover 13-21
Forcing Failover 13-21
Disabling Failover 13-21
Restoring a Failed Unit or Failover Group 13-22
Monitoring Active/Active Failover 13-22

System 13-22
Failover Group 1 and Failover Group 2 13-23
Feature History for Active/Active Failover 13-23

PART 2 Configuring Interfaces

CHAPTER 15 Starting Interface Configuration (ASA 5510 and Higher) 15-1
Information About Starting ASA 5510 and Higher Interface Configuration 15-2
Auto-MDI/MDIX Feature 15-2
Interfaces in Transparent Mode 15-2
Management Interface 15-2
Redundant Interfaces 15-5
EtherChannels 15-5
Licensing Requirements for ASA 5510 and Higher Interfaces 15-8
Guidelines and Limitations 15-10
Default Settings 15-12
Starting Interface Configuration (ASA 5510 and Higher) 15-12
Task Flow for Starting Interface Configuration 15-13
Converting In-Use Interfaces to a Redundant or EtherChannel Interface 15-14
Enabling the Physical Interface and Configuring Ethernet Parameters 15-23
Configuring a Redundant Interface 15-25
Configuring an EtherChannel 15-28
Configuring VLAN Subinterfaces and 802.1Q Trunking 15-34
Enabling Jumbo Frame Support (Supported Models) 15-37
Monitoring Interfaces 15-37
ARP Table 15-38
MAC Address Table 15-38
Interface Graphs 15-38
Where to Go Next 15-41
Feature History for ASA 5510 and Higher Interfaces 15-41

CHAPTER 16 Starting Interface Configuration (ASA 5505) 16-1
Information About ASA 5505 Interfaces 16-1
Understanding ASA 5505 Ports and Interfaces 16-2
Maximum Active VLAN Interfaces for Your License 16-2
VLAN MAC Addresses 16-4
Power over Ethernet 16-4
Monitoring Traffic Using SPAN 16-4

Auto-MDI/MDIX Feature 16-4
Licensing Requirements for ASA 5505 Interfaces 16-4
Guidelines and Limitations 16-5
Default Settings 16-5
Starting ASA 5505 Interface Configuration 16-6
Task Flow for Starting Interface Configuration 16-6
Configuring VLAN Interfaces 16-6
Configuring and Enabling Switch Ports as Access Ports 16-8
Configuring and Enabling Switch Ports as Trunk Ports 16-10
Monitoring Interfaces 16-12
ARP Table 16-12
MAC Address Table 16-12
Interface Graphs 16-13
Where to Go Next 16-15
Feature History for ASA 5505 Interfaces 16-16

CHAPTER 17 **Completing Interface Configuration (Routed Mode)** 17-1
Information About Completing Interface Configuration in Routed Mode 17-1
Security Levels 17-1
Dual IP Stack (IPv4 and IPv6) 17-2
Licensing Requirements for Completing Interface Configuration in Routed Mode 17-2
Guidelines and Limitations 17-5
Default Settings 17-6
Completing Interface Configuration in Routed Mode 17-6
Task Flow for Completing Interface Configuration 17-7
Configuring General Interface Parameters 17-7
Configuring the MAC Address and MTU 17-13
Configuring IPv6 Addressing 17-15
Allowing Same Security Level Communication 17-20
Turning Off and Turning On Interfaces 17-22
Monitoring Interfaces 17-22
ARP Table 17-23
DHCP 17-23
MAC Address Table 17-26
Dynamic ACLs 17-26
Interface Graphs 17-26
PPPoE Client 17-29
Interface Connection 17-29

Feature History for Interfaces in Routed Mode 17-30

CHAPTER 18 **Completing Interface Configuration (Transparent Mode, 8.4 and Later)** 18-1

Information About Completing Interface Configuration in Transparent Mode (8.4 and Later) 18-1
Bridge Groups in Transparent Mode 18-2
Security Levels 18-2
Licensing Requirements for Completing Interface Configuration in Transparent Mode 18-3
Guidelines and Limitations 18-5
Default Settings 18-6
Completing Interface Configuration in Transparent Mode (8.4 and Later) 18-7
Task Flow for Completing Interface Configuration 18-7
Configuring Bridge Groups 18-8
Configuring General Interface Parameters 18-9
Configuring a Management Interface (ASA 5510 and Higher) 18-11
Configuring the MAC Address and MTU 18-14
Configuring IPv6 Addressing 18-16
Allowing Same Security Level Communication 18-21
Turning Off and Turning On Interfaces 18-21
Monitoring Interfaces 18-22
ARP Table 18-22
DHCP 18-23
MAC Address Table 18-25
Dynamic ACLs 18-25
Interface Graphs 18-26
PPPoE Client 18-28
Interface Connection 18-28
Feature History for Interfaces in Transparent Mode 18-30

CHAPTER 19 **Completing Interface Configuration (Transparent Mode, 8.3 and Earlier)** 19-1

Information About Completing Interface Configuration in Transparent Mode (8.3 and Earlier) 19-1
Information About the Global Management IP Address 19-2
Security Levels 19-2
Licensing Requirements for Completing Interface Configuration in Transparent Mode 19-3
Guidelines and Limitations 19-3
Default Settings 19-4
Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier) 19-4
Configuring the IPv4 Address 19-4
Configuring the IPv6 Address 19-5

Completing Interface Configuration in Transparent Mode (8.3 and Earlier) 19-9
Task Flow for Completing Interface Configuration 19-9
Configuring General Interface Parameters 19-10
Configuring a Management Interface (ASA 5510 and Higher) 19-11
Configuring the MAC Address and MTU 19-15
Allowing Same Security Level Communication 19-17
Monitoring Interfaces 19-17
Feature History for Interfaces in Transparent Mode 19-18

PART 3 **Configuring Basic Settings**

CHAPTER 21 **Configuring Basic Settings** 21-1
Configuring the Hostname, Domain Name, and Passwords 21-1
Setting the Hostname, Domain Name, and the enable and Telnet Passwords 21-1
Setting the Date and Time 21-2
Setting the Date and Time Using an NTP Server 21-3
Setting the Date and Time Manually 21-4
Configuring the Master Passphrase 21-4
Information About the Master Passphrase 21-5
Licensing Requirements for the Master Passphrase 21-5
Guidelines and Limitations 21-5
Adding or Changing the Master Passphrase 21-5
Disabling the Master Passphrase 21-6
Recovering the Master Passphrase 21-7
Feature History for the Master Passphrase 21-7
Configuring the DNS Server 21-8
Changing the Heap Memory Size 21-9
Monitoring DNS Cache 21-9

CHAPTER 22 **Configuring Dynamic DNS** 22-1
Information About DDNS 22-1
Licensing Requirements for DDNS 22-2
Guidelines and Limitations 22-2
Configuring Dynamic DNS 22-2
DDNS Monitoring 22-4
Feature History for DDNS 22-4

CHAPTER 23 **Configuring DHCP** 23-1

Information About DHCP 23-1

Licensing Requirements for DHCP 23-2

Guidelines and Limitations 23-2

Configuring a DHCP Server 23-3

Editing DHCP Server Settings 23-4

Configuring Advanced DHCP Options 23-5

Configuring the DHCP Relay Service 23-6

Editing DHCP Relay Agent Settings 23-7

Adding or Editing Global DHCP Relay Server Settings 23-8

Additional References 23-8

RFCs 23-9

DHCP Monitoring 23-9

Feature History for DHCP 23-10

PART 4 **Configuring Objects and ACLs**

CHAPTER 25 **Configuring Objects** 25-1

Information About Objects 25-1

Licensing Requirements for Objects 25-1

Guidelines and Limitations 25-1

Configuring Objects 25-2

Configuring Network Objects and Groups 25-2

Configuring Service Objects and Service Groups 25-4

Configuring Local User Groups 25-7

Configuring Security Group Object Groups 25-8

Configuring Regular Expressions 25-10

Configuring Time Ranges 25-15

Monitoring Objects 25-16

Feature History for Objects 25-16

CHAPTER 26 **Using the ACL Manager** 26-1

Information About the ACL Manager 26-1

Licensing Requirements for the ACL Manager 26-1

Guidelines and Limitations 26-2

Adding ACLs and ACEs 26-2

Using Standard ACLs in the ACL Manager 26-4

Feature History for the ACL Manager 26-5

CHAPTER 27 **Adding a Standard Access Control List** 27-1

Information About Standard ACLs 27-1

Licensing Requirements for Standard ACLs 27-1

Guidelines and Limitations 27-1

Default Settings 27-2

Adding Standard ACLs 27-3

Using Standard ACLs 27-3

Feature History for Standard ACLs 27-4

CHAPTER 28 **Adding a Webtype Access Control List** 28-1

Licensing Requirements for Webtype ACLs 28-1

Guidelines and Limitations 28-1

Default Settings 28-2

Using Webtype ACLs 28-2

Task Flow for Configuring Webtype ACLs 28-2

Adding a Webtype ACL and ACE 28-3

Editing Webtype ACLs and ACEs 28-4

Deleting Webtype ACLs and ACEs 28-5

Feature History for Webtype Access Lists 28-6

PART 5 **Configuring IP Routing**

CHAPTER 30 **Routing Overview** 30-1

Information About Routing 30-1

Switching 30-1

Path Determination 30-2

Supported Route Types 30-2

How Routing Behaves Within the ASA 30-4

Egress Interface Selection Process 30-4

Next Hop Selection Process 30-4

Supported Internet Protocols for Routing 30-5

Information About the Routing Table 30-5

Displaying the Routing Table 30-6

How the Routing Table Is Populated 30-6

How Forwarding Decisions Are Made 30-8

Dynamic Routing and Failover 30-8
Dynamic Routing and Clustering 30-9
Dynamic Routing in Multiple Context Mode 30-10
Disabling Proxy ARPs 30-11

CHAPTER 31 **Configuring Static and Default Routes** 31-1
Information About Static and Default Routes 31-1
Licensing Requirements for Static and Default Routes 31-2
Guidelines and Limitations 31-2
Configuring Static and Default Routes 31-2
Configuring a Static Route 31-3
Configuring a Default Static Route 31-6
Configuring IPv6 Default and Static Routes 31-7
Monitoring a Static or Default Route 31-8
Configuration Examples for Static or Default Routes 31-8
Feature History for Static and Default Routes 31-9

CHAPTER 32 **Defining Route Maps** 32-1
Information About Route Maps 32-1
Permit and Deny Clauses 32-2
Match and Set Clause Values 32-2
Licensing Requirements for Route Maps 32-3
Guidelines and Limitations 32-3
Defining a Route Map 32-4
Customizing a Route Map 32-5
Defining a Route to Match a Specific Destination Address 32-5
Configuring Prefix Lists 32-6
Configuring Prefix Rules 32-7
Configuring the Metric Values for a Route Action 32-7
Configuration Example for Route Maps 32-8
Feature History for Route Maps 32-8

CHAPTER 33 **Configuring OSPF** 33-1
Information About OSPF 33-1
Implementation Differences Between OSPFv2 and OSPFv3 33-3
Using Clustering 33-3
Licensing Requirements for OSPF 33-3

Guidelines and Limitations 33-3
Configuring OSPFv2 33-5
Customizing OSPFv2 33-6
Redistributing Routes Into OSPFv2 33-6
Configuring Route Summarization When Redistributing Routes Into OSPFv2 33-8
Configuring Route Summarization Between OSPFv2 Areas 33-10
Configuring OSPFv2 Interface Parameters 33-11
Configuring OSPFv2 Area Parameters 33-13
Configuring an OSPFv2 NSSA 33-14
Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3) 33-15
Defining Static OSPFv2 Neighbors 33-17
Configuring Route Calculation Timers 33-18
Logging Neighbors Going Up or Down 33-18
Configuring Filtering in OSPF 33-19
Configuring a Virtual Link in OSPF 33-20
Configuring OSPFv3 33-21
Enabling OSPFv3 33-22
Configuring OSPFv3 Interface Parameters 33-22
Configuring OSPFv3 Area Parameters 33-23
Configuring a Virtual Link Neighbor 33-24
Configuring OSPFv3 Passive Interfaces 33-25
Configuring OSPFv3 Administrative Distance 33-26
Configuring OSPFv3 Timers 33-27
Defining Static OSPFv3 Neighbors 33-28
Sending Syslog Messages 33-29
Suppressing Syslog Messages 33-29
Calculating Summary Route Costs 33-29
Generating a Default External Route into an OSPFv3 Routing Domain 33-30
Configuring an IPv6 Summary Prefix 33-30
Redistributing IPv6 Routes 33-31
Removing the OSPF Configuration 33-32
Configuration Example for OSPFv2 33-32
Configuration Example for OSPFv3 33-34
Monitoring OSPF 33-35
Additional References 33-36
RFCs 33-36
Feature History for OSPF 33-37

CHAPTER 34 **Configuring RIP** 34-1

Information About RIP 34-1
Routing Update Process 34-2
RIP Routing Metric 34-2
RIP Stability Features 34-2
RIP Timers 34-2
Using Clustering 34-3
Licensing Requirements for RIP 34-3
Guidelines and Limitations 34-3
Configuring RIP 34-4
Enabling RIP 34-4
Customizing RIP 34-5
Configuring the RIP Version 34-5
Configuring Interfaces for RIP 34-5
Configuring the RIP Send and Receive Version on an Interface 34-7
Configuring Route Summarization 34-8
Filtering Networks in RIP 34-8
Redistributing Routes into the RIP Routing Process 34-10
Enabling RIP Authentication 34-11
Restarting the RIP Process 34-12
Monitoring RIP 34-12
Configuration Example for RIP 34-12
Feature History for RIP 34-13

CHAPTER 35 **Configuring Multicast Routing** 35-1

Information About Multicast Routing 35-1
Stub Multicast Routing 35-2
PIM Multicast Routing 35-2
Multicast Group Concept 35-2
Clustering 35-2
Licensing Requirements for Multicast Routing 35-3
Guidelines and Limitations 35-3
Enabling Multicast Routing 35-3
Customizing Multicast Routing 35-4
Configuring Stub Multicast Routing and Forwarding IGMP Messages 35-4
Configuring a Static Multicast Route 35-5
Configuring IGMP Features 35-6
Configuring PIM Features 35-10

Configuring a Multicast Group 35-14
Configuring a Bidirectional Neighbor Filter 35-15
Configuring a Multicast Boundary 35-16
Configuration Example for Multicast Routing 35-17
Additional References 35-18
Related Documents 35-19
RFCs 35-19
Feature History for Multicast Routing 35-19

CHAPTER 36 **Configuring EIGRP** 36-1
Information About EIGRP 36-1
Using Clustering 36-2
Licensing Requirements for EIGRP 36-2
Guidelines and Limitations 36-3
Task List to Configure an EIGRP Process 36-3
Configuring EIGRP 36-4
Enabling EIGRP 36-4
Enabling EIGRP Stub Routing 36-5
Customizing EIGRP 36-6
Defining a Network for an EIGRP Routing Process 36-7
Configuring Interfaces for EIGRP 36-7
Configuring the Summary Aggregate Addresses on Interfaces 36-9
Changing the Interface Delay Value 36-10
Enabling EIGRP Authentication on an Interface 36-10
Defining an EIGRP Neighbor 36-11
Redistributing Routes Into EIGRP 36-12
Filtering Networks in EIGRP 36-13
Customizing the EIGRP Hello Interval and Hold Time 36-14
Disabling Automatic Route Summarization 36-15
Configuring Default Information in EIGRP 36-16
Disabling EIGRP Split Horizon 36-17
Restarting the EIGRP Process 36-17
Monitoring EIGRP 36-18
Feature History for EIGRP 36-19

CHAPTER 37 **Configuring IPv6 Neighbor Discovery** 37-1
Information About IPv6 Neighbor Discovery 37-1
Neighbor Solicitation Messages 37-2

Neighbor Reachable Time 37-3
Duplicate Address Detection 37-3
Router Advertisement Messages 37-3
Static IPv6 Neighbors 37-5
Licensing Requirements for IPv6 Neighbor Discovery 37-5
Prerequisites for IPv6 Neighbor Discovery 37-5
Guidelines and Limitations 37-5
Default Settings for IPv6 Neighbor Discovery 37-7
Configuring IPv6 Neighbor Discovery 37-7
Configuring the Neighbor Solicitation Message Interval 37-8
Configuring the Neighbor Reachable Time 37-8
Configuring the Router Advertisement Transmission Interval 37-9
Configuring the Router Lifetime Value 37-9
Configuring DAD Settings 37-10
Suppressing Router Advertisement Messages 37-10
Configuring Address Config Flags for IPv6 DHCP Relay 37-11
Configuring the IPv6 Prefix in Router Advertisements 37-11
Configuring a Static IPv6 Neighbor 37-12
Viewing and Clearing Dynamically Discovered Neighbors 37-13
Additional References 37-13
Related Documents for IPv6 Prefixes 37-14
RFCs for IPv6 Prefixes and Documentation 37-14
Feature History for IPv6 Neighbor Discovery 37-14

PART 6 **Configuring Network Address Translation**

CHAPTER 39 **Information About NAT (ASA 8.3 and Later)** 39-1
Why Use NAT? 39-1
NAT Terminology 39-2
NAT Types 39-3
NAT Types Overview 39-3
Static NAT 39-3
Dynamic NAT 39-8
Dynamic PAT 39-10
Identity NAT 39-12
NAT in Routed and Transparent Mode 39-12
NAT in Routed Mode 39-13
NAT in Transparent Mode 39-13

NAT and IPv6 39-15
How NAT is Implemented 39-15
Main Differences Between Network Object NAT and Twice NAT 39-15
Information About Network Object NAT 39-16
Information About Twice NAT 39-16
NAT Rule Order 39-20
NAT Interfaces 39-21
Routing NAT Packets 39-21
Mapped Addresses and Routing 39-22
Transparent Mode Routing Requirements for Remote Networks 39-24
Determining the Egress Interface 39-24
NAT for VPN 39-24
NAT and Remote Access VPN 39-24
NAT and Site-to-Site VPN 39-26
NAT and VPN Management Access 39-28
Troubleshooting NAT and VPN 39-30
DNS and NAT 39-30
Where to Go Next 39-34

CHAPTER 40 **Configuring Network Object NAT (ASA 8.3 and Later)** 40-1
Information About Network Object NAT 40-1
Licensing Requirements for Network Object NAT 40-2
Prerequisites for Network Object NAT 40-2
Guidelines and Limitations 40-2
Default Settings 40-3
Configuring Network Object NAT 40-4
Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool 40-4
Configuring Dynamic PAT (Hide) 40-8
Configuring Static NAT or Static NAT-with-Port-Translation 40-11
Configuring Identity NAT 40-15
Configuring Per-Session PAT Rules 40-18
Monitoring Network Object NAT 40-19
Configuration Examples for Network Object NAT 40-20
Providing Access to an Inside Web Server (Static NAT) 40-21
NAT for Inside Hosts (Dynamic NAT) and NAT for an Outside Web Server (Static NAT) 40-23
Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many) 40-28
Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation) 40-32

DNS Server on Mapped Interface, Web Server on Real Interface (Static NAT with DNS Modification) 40-35
DNS Server and FTP Server on Mapped Interface, FTP Server is Translated (Static NAT with DNS Modification) 40-38
IPv4 DNS Server and FTP Server on Mapped Interface, IPv6 Host on Real Interface (Static NAT64 with DNS64 Modification) 40-40
Feature History for Network Object NAT 40-45

CHAPTER 41 **Configuring Twice NAT (ASA 8.3 and Later)** 41-1
Information About Twice NAT 41-1
Licensing Requirements for Twice NAT 41-2
Prerequisites for Twice NAT 41-2
Guidelines and Limitations 41-2
Default Settings 41-4
Configuring Twice NAT 41-4
Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool 41-4
Configuring Dynamic PAT (Hide) 41-12
Configuring Static NAT or Static NAT-with-Port-Translation 41-18
Configuring Identity NAT 41-24
Configuring Per-Session PAT Rules 41-29
Monitoring Twice NAT 41-30
Configuration Examples for Twice NAT 41-30
Different Translation Depending on the Destination (Dynamic PAT) 41-31
Different Translation Depending on the Destination Address and Port (Dynamic PAT) 41-39
Feature History for Twice NAT 41-48

CHAPTER 42 **Configuring NAT (ASA 8.2 and Earlier)** 42-1
NAT Overview 42-1
Introduction to NAT 42-1
NAT in Routed Mode 42-2
NAT in Transparent Mode 42-3
NAT Control 42-4
NAT Types 42-6
Policy NAT 42-10
NAT and Same Security Level Interfaces 42-12
Order of NAT Rules Used to Match Real Addresses 42-13
Mapped Address Guidelines 42-13
DNS and NAT 42-13
Configuring NAT Control 42-15

Using Dynamic NAT 42-16
Dynamic NAT Implementation 42-16
Managing Global Pools 42-21
Configuring Dynamic NAT, PAT, or Identity NAT 42-22
Configuring Dynamic Policy NAT or PAT 42-24
Using Static NAT 42-26
Configuring Static NAT, PAT, or Identity NAT 42-27
Configuring Static Policy NAT, PAT, or Identity NAT 42-30
Using NAT Exemption 42-32

PART 7 Configuring Service Policies

CHAPTER 44 Configuring a Service Policy 44-1
Information About Service Policies 44-1
Supported Features 44-1
Feature Directionality 44-2
Feature Matching Within a Service Policy 44-3
Order in Which Multiple Feature Actions are Applied 44-4
Incompatibility of Certain Feature Actions 44-4
Feature Matching for Multiple Service Policies 44-5
Licensing Requirements for Service Policies 44-5
Guidelines and Limitations 44-6
Default Settings 44-7
Default Configuration 44-7
Default Traffic Classes 44-8
Task Flows for Configuring Service Policies 44-8
Task Flow for Configuring a Service Policy Rule 44-8
Adding a Service Policy Rule for Through Traffic 44-8
Adding a Service Policy Rule for Management Traffic 44-13
Configuring a Service Policy Rule for Management Traffic 44-13
Managing the Order of Service Policy Rules 44-16
Feature History for Service Policies 44-17

PART 8 Configuring AAA Servers and the Local Database

CHAPTER 46 Configuring AAA Servers and the Local Database 46-1
Information About AAA 46-1
Information About Authentication 46-2

Information About Authorization 46-2
Information About Accounting 46-3
Summary of Server Support 46-3
RADIUS Server Support 46-4
TACACS+ Server Support 46-5
RSA/SDI Server Support 46-5
NT Server Support 46-6
Kerberos Server Support 46-6
LDAP Server Support 46-6
HTTP Forms Authentication for Clientless SSL VPN 46-7
Local Database Support, Including as a Fallback Method 46-7
How Fallback Works with Multiple Servers in a Group 46-8
Using Certificates and User Login Credentials 46-8
Licensing Requirements for AAA Servers 46-9
Guidelines and Limitations 46-10
Configuring AAA 46-10
Task Flow for Configuring AAA 46-10
Configuring AAA Server Groups 46-11
Adding a Server to a Group 46-12
Configuring AAA Server Parameters 46-13
Configuring LDAP Attribute Maps 46-19
Adding a User Account to the Local Database 46-20
Authenticating Users with a Public Key for SSH 46-25
Adding an Authentication Prompt 46-26
Testing Server Authentication and Authorization 46-27
Monitoring AAA Servers 46-27
Additional References 46-28
RFCs 46-28
Feature History for AAA Servers 46-28

CHAPTER 47 **Configuring the Identity Firewall** 47-1
Information About the Identity Firewall 47-1
Overview of the Identity Firewall 47-1
Architecture for Identity Firewall Deployments 47-2
Features of the Identity Firewall 47-3
Deployment Scenarios 47-4
Licensing for the Identity Firewall 47-7
Guidelines and Limitations 47-7

Prerequisites 47-8
Configuring the Identity Firewall 47-9
Task Flow for Configuring the Identity Firewall 47-9
Configuring the Active Directory Domain 47-10
Configuring Active Directory Server Groups 47-10
Configuring Active Directory Agents 47-11
Configuring Active Directory Agent Groups 47-12
Configuring Identity Options 47-12
Configuring Identity-Based Security Policy 47-15
Monitoring the Identity Firewall 47-16
Monitoring AD Agents 47-16
Monitoring Groups 47-16
Monitoring Memory Usage for the Identity Firewall 47-17
Monitoring Users for the Identity Firewall 47-17
Feature History for the Identity Firewall 47-18

CHAPTER 48 **Configuring the ASA to Integrate with Cisco TrustSec** 48-1
Information About the ASA Integrated with Cisco TrustSec 48-1
Information about Cisco TrustSec 48-1
About SGT and SXP Support in Cisco TrustSec 48-2
Roles in the Cisco TrustSec Solution 48-3
Security Group Policy Enforcement 48-3
How the ASA Enforces Security Group Based Policies 48-4
About Speaker and Listener Roles on the ASA 48-6
Features of the ASA-Cisco TrustSec Integration 48-6
Licensing Requirements when Integrating the ASA with Cisco TrustSec 48-8
Prerequisites for Integrating the ASA with Cisco TrustSec 48-8
Guidelines and Limitations 48-9
Configuring the ASA for Cisco TrustSec Integration 48-11
Task Flow for Configuring the ASA to Integrate with Cisco TrustSec 48-11
Configuring the AAA Server for Cisco TrustSec Integration 48-12
Importing a Protected Access Credential (PAC) File 48-13
Configuring the Security Exchange Protocol (SXP) 48-14
Adding an SXP Connection Peer 48-15
Refreshing Environment Data 48-16
Configuring the Security Policy 48-17
Monitoring the ASA Integrated with Cisco TrustSec 48-17
Monitoring SXP Connections 48-17

Monitoring Environment Data 48-18
Monitoring Cisco TrustSec IP-SGT Mappings 48-19
Monitoring the PAC File 48-20
Feature History for the ASA-Cisco TrustSec Integration 48-21

CHAPTER 49 **Configuring Digital Certificates** 49-1
Information About Digital Certificates 49-1
Public Key Cryptography 49-3
Certificate Scalability 49-3
Key Pairs 49-4
Trustpoints 49-4
Revocation Checking 49-5
The Local CA 49-7
Licensing Requirements for Digital Certificates 49-8
Prerequisites for Local Certificates 49-8
Prerequisites for SCEP Proxy Support 49-9
Guidelines and Limitations 49-9
Configuring Digital Certificates 49-10
Configuring CA Certificate Authentication 49-11
Adding or Installing a CA Certificate 49-11
Editing or Removing a CA Certificate Configuration 49-12
Showing CA Certificate Details 49-13
Configuring CA Certificate Authentication 49-13
Adding or Installing a CA Certificate 49-13
Editing or Removing a CA Certificate Configuration 49-14
Showing CA Certificate Details 49-14
Configuring CA Certificates for Revocation 49-15
Configuring CRL Retrieval Policy 49-15
Configuring CRL Retrieval Methods 49-16
Configuring OCSP Rules 49-16
Configuring Advanced CRL and OCSP Settings 49-17
Monitoring CRLs 49-18
Configuring CA Certificates for Revocation 49-19
Configuring CRL Retrieval Policy 49-19
Configuring CRL Retrieval Methods 49-20
Configuring OCSP Rules 49-20
Configuring Advanced CRL and OCSP Settings 49-21
Configuring Identity Certificates Authentication 49-22
Adding or Importing an Identity Certificate 49-22

Showing Identity Certificate Details 49-24
Deleting an Identity Certificate 49-25
Exporting an Identity Certificate 49-25
Generating a Certificate Signing Request 49-25
Installing Identity Certificates 49-26
Configuring Code Signer Certificates 49-27
Showing Code Signer Certificate Details 49-28
Deleting a Code Signer Certificate 49-28
Importing a Code Signer Certificate 49-28
Exporting a Code Signer Certificate 49-29
Authenticating Using the Local CA 49-29
Configuring the Local CA Server 49-30
Deleting the Local CA Server 49-32
Managing the User Database 49-33
Adding a Local CA User 49-33
Sending an Initial OTP or Replacing OTPs 49-34
Editing a Local CA User 49-34
Deleting a Local CA User 49-35
Allowing User Enrollment 49-35
Viewing or Regenerating an OTP 49-35
Managing User Certificates 49-35
Monitoring CRLs 49-36
Feature History for Certificate Management 49-37

PART 9 **Configuring Access Control**

CHAPTER 51 **Configuring Access Rules** 51-1
Information About Access Rules 51-1
General Information About Rules 51-2
Information About Access Rules 51-4
Information About EtherType Rules 51-5
Licensing Requirements for Access Rules 51-6
Guidelines and Limitations 51-6
Default Settings 51-7
Configuring Access Rules 51-7
Adding an Access Rule 51-7
Adding an EtherType Rule (Transparent Mode Only) 51-8
Configuring Management Access Rules 51-9

Advanced Access Rule Configuration 51-10
Configuring HTTP Redirect 51-12
Feature History for Access Rules 51-13

CHAPTER 52 **Configuring Management Access** 52-1
Configuring ASA Access for ASDM, Telnet, or SSH 52-1
Licensing Requirements for ASA Access for ASDM, Telnet, or SSH 52-1
Guidelines and Limitations 52-2
Configuring Management Access 52-3
Using a Telnet Client 52-4
Using an SSH Client 52-4
Configuring CLI Parameters 52-5
Licensing Requirements for CLI Parameters 52-5
Guidelines and Limitations 52-5
Configuring a Login Banner 52-5
Customizing a CLI Prompt 52-6
Changing the Console Timeout 52-8
Configuring File Access 52-8
Licensing Requirements for File Access 52-8
Guidelines and Limitations 52-8
Configuring the FTP Client Mode 52-9
Configuring the ASA as a Secure Copy Server 52-9
Configuring the ASA as a TFTP Client 52-10
Adding Mount Points 52-10
Configuring ICMP Access 52-12
Information About ICMP Access 52-12
Licensing Requirements for ICMP Access 52-12
Guidelines and Limitations 52-13
Default Settings 52-13
Configuring ICMP Access 52-13
Configuring Management Access Over a VPN Tunnel 52-14
Licensing Requirements for a Management Interface 52-14
Guidelines and Limitations 52-14
Configuring a Management Interface 52-15
Configuring AAA for System Administrators 52-15
Information About AAA for System Administrators 52-16
Licensing Requirements for AAA for System Administrators 52-19
Prerequisites 52-19
Guidelines and Limitations 52-20

Default Settings 52-20
Configuring Authentication for CLI, ASDM, and enable command Access 52-20
Limiting User CLI and ASDM Access with Management Authorization 52-21
Configuring Command Authorization 52-23
Configuring Management Access Accounting 52-29
Viewing the Currently Logged-In User 52-29
Recovering from a Lockout 52-30
Monitoring Device Access 52-31
Feature History for Management Access 52-32

CHAPTER 53 **Configuring AAA Rules for Network Access** 53-1
AAA Performance 53-1
Licensing Requirements for AAA Rules 53-1
Guidelines and Limitations 53-2
Configuring Authentication for Network Access 53-2
Information About Authentication 53-2
Configuring Network Access Authentication 53-6
Enabling the Redirection Method of Authentication for HTTP and HTTPS 53-7
Enabling Secure Authentication of Web Clients 53-8
Authenticating Directly with the ASA 53-9
Configuring the Authentication Proxy Limit 53-11
Configuring Authorization for Network Access 53-11
Configuring TACACS+ Authorization 53-11
Configuring RADIUS Authorization 53-13
Configuring Accounting for Network Access 53-17
Using MAC Addresses to Exempt Traffic from Authentication and Authorization 53-18
Feature History for AAA Rules 53-19

CHAPTER 54 **Configuring Web Cache Services Using WCCP** 54-1
Information About WCCP 54-1
Guidelines and Limitations 54-1
Licensing Requirements for WCCP 54-3
Configuring WCCP Service Groups 54-3
Adding or Editing WCCP Service Groups 54-3
Configuring Packet Redirection 54-4
Adding or Editing Packet Redirection 54-4
WCCP Monitoring 54-4

Feature History for WCCP 54-5

CHAPTER 55 **Configuring Public Servers** 55-1

Information About Public Servers 55-1

Licensing Requirements for Public Servers 55-1

Guidelines and Limitations 55-1

Adding a Public Server that Enables Static NAT 55-2

Adding a Public Server that Enables Static NAT with PAT 55-2

Editing Settings for a Public Server 55-3

Feature History for Public Servers 55-4

PART 10 **Configuring Application Inspection**

CHAPTER 57 **Getting Started with Application Layer Protocol Inspection** 57-1

Information about Application Layer Protocol Inspection 57-1

How Inspection Engines Work 57-1

When to Use Application Protocol Inspection 57-2

Guidelines and Limitations 57-3

Default Settings 57-4

Configuring Application Layer Protocol Inspection 57-6

CHAPTER 58 **Configuring Inspection of Basic Internet Protocols** 58-1

DNS Inspection 58-1

How DNS Application Inspection Works 58-2

How DNS Rewrite Works 58-3

Configuring DNS Rewrite 58-3

Select DNS Inspect Map 58-5

DNS Class Map 58-6

Add/Edit DNS Traffic Class Map 58-6

Add/Edit DNS Match Criterion 58-7

DNS Inspect Map 58-8

Add/Edit DNS Policy Map (Security Level) 58-9

Add/Edit DNS Policy Map (Details) 58-11

FTP Inspection 58-12

FTP Inspection Overview 58-12

Using Strict FTP 58-13

Select FTP Map 58-14

FTP Class Map 58-14

Add/Edit FTP Traffic Class Map 58-15
Add/Edit FTP Match Criterion 58-15
FTP Inspect Map 58-16
File Type Filtering 58-17
Add/Edit FTP Policy Map (Security Level) 58-18
Add/Edit FTP Policy Map (Details) 58-18
Add/Edit FTP Map 58-19
Verifying and Monitoring FTP Inspection 58-21
HTTP Inspection 58-21
HTTP Inspection Overview 58-21
Select HTTP Map 58-22
HTTP Class Map 58-22
Add/Edit HTTP Traffic Class Map 58-23
Add/Edit HTTP Match Criterion 58-23
HTTP Inspect Map 58-27
URI Filtering 58-28
Add/Edit HTTP Policy Map (Security Level) 58-28
Add/Edit HTTP Policy Map (Details) 58-29
Add/Edit HTTP Map 58-30
ICMP Inspection 58-34
ICMP Error Inspection 58-34
Instant Messaging Inspection 58-35
IM Inspection Overview 58-35
Adding a Class Map for IM Inspection 58-35
Select IM Map 58-36
IP Options Inspection 58-36
IP Options Inspection Overview 58-37
Configuring IP Options Inspection 58-37
Select IP Options Inspect Map 58-39
IP Options Inspect Map 58-39
Add/Edit IP Options Inspect Map 58-39
IPsec Pass Through Inspection 58-40
IPsec Pass Through Inspection Overview 58-40
Select IPsec-Pass-Thru Map 58-41
IPsec Pass Through Inspect Map 58-41
Add/Edit IPsec Pass Thru Policy Map (Security Level) 58-42
Add/Edit IPsec Pass Thru Policy Map (Details) 58-42
IPv6 Inspection 58-43
Configuring an IPv6 Inspection Policy Map 58-44

NetBIOS Inspection 58-45
NetBIOS Inspection Overview 58-45
Select NETBIOS Map 58-45
NetBIOS Inspect Map 58-45
Add/Edit NetBIOS Policy Map 58-46
PPTP Inspection 58-46
SMTP and Extended SMTP Inspection 58-47
SMTP and ESMTP Inspection Overview 58-47
Select ESMTP Map 58-48
ESMTP Inspect Map 58-48
MIME File Type Filtering 58-49
Add/Edit ESMTP Policy Map (Security Level) 58-50
Add/Edit ESMTP Policy Map (Details) 58-51
Add/Edit ESMTP Inspect 58-52
TFTP Inspection 58-55

CHAPTER 59 **Configuring Inspection for Voice and Video Protocols** 59-1

CTIQBE Inspection 59-1
CTIQBE Inspection Overview 59-1
Limitations and Restrictions 59-2
H.323 Inspection 59-2
H.323 Inspection Overview 59-3
How H.323 Works 59-3
H.239 Support in H.245 Messages 59-4
Limitations and Restrictions 59-4
Select H.323 Map 59-5
H.323 Class Map 59-5
Add/Edit H.323 Traffic Class Map 59-6
Add/Edit H.323 Match Criterion 59-6
H.323 Inspect Map 59-7
Phone Number Filtering 59-8
Add/Edit H.323 Policy Map (Security Level) 59-8
Add/Edit H.323 Policy Map (Details) 59-9
Add/Edit HSI Group 59-11
Add/Edit H.323 Map 59-11
MGCP Inspection 59-12
MGCP Inspection Overview 59-12
Select MGCP Map 59-14
MGCP Inspect Map 59-14

Gateways and Call Agents 59-15
Add/Edit MGCP Policy Map 59-15
Add/Edit MGCP Group 59-16
RTSP Inspection 59-16
RTSP Inspection Overview 59-17
Using RealPlayer 59-17
Restrictions and Limitations 59-18
Select RTSP Map 59-18
RTSP Inspect Map 59-18
Add/Edit RTSP Policy Map 59-18
Add/Edit RTSP Inspect 59-19
SIP Inspection 59-20
SIP Inspection Overview 59-20
SIP Instant Messaging 59-21
Select SIP Map 59-22
SIP Class Map 59-23
Add/Edit SIP Traffic Class Map 59-23
Add/Edit SIP Match Criterion 59-24
SIP Inspect Map 59-26
Add/Edit SIP Policy Map (Security Level) 59-27
Add/Edit SIP Policy Map (Details) 59-28
Add/Edit SIP Inspect 59-29
Skinny (SCCP) Inspection 59-31
SCCP Inspection Overview 59-32
Supporting Cisco IP Phones 59-32
Restrictions and Limitations 59-33
Select SCCP (Skinny) Map 59-33
SCCP (Skinny) Inspect Map 59-34
Message ID Filtering 59-35
Add/Edit SCCP (Skinny) Policy Map (Security Level) 59-35
Add/Edit SCCP (Skinny) Policy Map (Details) 59-36
Add/Edit Message ID Filter 59-37

CHAPTER 60 **Configuring Inspection of Database and Directory Protocols** 60-1
ILS Inspection 60-1
SQL*Net Inspection 60-2
Sun RPC Inspection 60-3
Sun RPC Inspection Overview 60-3
SUNRPC Server 60-4

Add/Edit SUNRPC Service 60-4

CHAPTER 61 **Configuring Inspection for Management Application Protocols** 61-1

DCERPC Inspection 61-1
DCERPC Overview 61-1
Select DCERPC Map 61-2
DCERPC Inspect Map 61-2
Add/Edit DCERPC Policy Map 61-3
GTP Inspection 61-4
GTP Inspection Overview 61-5
Select GTP Map 61-5
GTP Inspect Map 61-6
IMSI Prefix Filtering 61-7
Add/Edit GTP Policy Map (Security Level) 61-7
Add/Edit GTP Policy Map (Details) 61-8
Add/Edit GTP Map 61-9
RADIUS Accounting Inspection 61-10
RADIUS Accounting Inspection Overview 61-11
Select RADIUS Accounting Map 61-11
Add RADIUS Accounting Policy Map 61-11
RADIUS Inspect Map 61-12
RADIUS Inspect Map Host 61-12
RADIUS Inspect Map Other 61-13
RSH Inspection 61-13
SNMP Inspection 61-13
SNMP Inspection Overview 61-14
Select SNMP Map 61-14
SNMP Inspect Map 61-14
XDMCP Inspection 61-15

PART 11 **Configuring Unified Communications**

CHAPTER 63 **Information About Cisco Unified Communications Proxy Features** 63-1

Information About the Adaptive Security Appliance in Cisco Unified Communications 63-1
TLS Proxy Applications in Cisco Unified Communications 63-3
Licensing for Cisco Unified Communications Proxy Features 63-4

CHAPTER 64 **Using the Cisco Unified Communication Wizard** 64-1

Information about the Cisco Unified Communication Wizard 64-1

Licensing Requirements for the Unified Communication Wizard 64-3

Guidelines and Limitations 64-4

Configuring the Phone Proxy by using the Unified Communication Wizard 64-4

Configuring the Private Network for the Phone Proxy 64-5

Configuring Servers for the Phone Proxy 64-6

Enabling Certificate Authority Proxy Function (CAPF) for IP Phones 64-8

Configuring the Public IP Phone Network 64-9

Configuring the Media Termination Address for Unified Communication Proxies 64-10

Configuring the Mobility Advantage by using the Unified Communication Wizard 64-11

Configuring the Topology for the Cisco Mobility Advantage Proxy 64-12

Configuring the Server-Side Certificates for the Cisco Mobility Advantage Proxy 64-12

Configuring the Client-Side Certificates for the Cisco Mobility Advantage Proxy 64-13

Configuring the Presence Federation Proxy by using the Unified Communication Wizard 64-14

Configuring the Topology for the Cisco Presence Federation Proxy 64-14

Configuring the Local-Side Certificates for the Cisco Presence Federation Proxy 64-15

Configuring the Remote-Side Certificates for the Cisco Presence Federation Proxy 64-15

Configuring the UC-IME by using the Unified Communication Wizard 64-16

Configuring the Topology for the Cisco Intercompany Media Engine Proxy 64-17

Configuring the Private Network Settings for the Cisco Intercompany Media Engine Proxy 64-18

Adding a Cisco Unified Communications Manager Server for the UC-IME Proxy 64-20

Configuring the Public Network Settings for the Cisco Intercompany Media Engine Proxy 64-20

Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy 64-21

Configuring the Remote-Side Certificates for the Cisco Intercompany Media Engine Proxy 64-22

Working with Certificates in the Unified Communication Wizard 64-23

Exporting an Identity Certificate 64-23

Installing a Certificate 64-23

Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy 64-24

Saving the Identity Certificate Request 64-25

Installing the ASA Identity Certificate on the Mobility Advantage Server 64-26

Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers 64-26

CHAPTER 65 **Configuring the Cisco Phone Proxy** 65-1

Information About the Cisco Phone Proxy 65-1

Phone Proxy Functionality 65-1

Supported Cisco UCM and IP Phones for the Phone Proxy 65-3

Licensing Requirements for the Phone Proxy 65-4
Prerequisites for the Phone Proxy 65-6
Media Termination Instance Prerequisites 65-6
Certificates from the Cisco UCM 65-7
DNS Lookup Prerequisites 65-7
Cisco Unified Communications Manager Prerequisites 65-7
Access List Rules 65-7
NAT and PAT Prerequisites 65-8
Prerequisites for IP Phones on Multiple Interfaces 65-8
7960 and 7940 IP Phones Support 65-9
Cisco IP Communicator Prerequisites 65-10
Prerequisites for Rate Limiting TFTP Requests 65-10
End-User Phone Provisioning 65-11
Phone Proxy Guidelines and Limitations 65-12
Configuring the Phone Proxy 65-14
Task Flow for Configuring the Phone Proxy 65-14
Creating the CTL File 65-14
Adding or Editing a Record Entry in a CTL File 65-16
Creating the Media Termination Instance 65-17
Creating the Phone Proxy Instance 65-18
Adding or Editing the TFTP Server for a Phone Proxy 65-20
Configuring Linksys Routers with UDP Port Forwarding for the Phone Proxy 65-21
Feature History for the Phone Proxy 65-22

CHAPTER 66 **Configuring the TLS Proxy for Encrypted Voice Inspection** 66-1
Information about the TLS Proxy for Encrypted Voice Inspection 66-1
Decryption and Inspection of Unified Communications Encrypted Signaling 66-2
Supported Cisco UCM and IP Phones for the TLS Proxy 66-3
Licensing for the TLS Proxy 66-4
Prerequisites for the TLS Proxy for Encrypted Voice Inspection 66-6
Configuring the TLS Proxy for Encrypted Voice Inspection 66-6
CTL Provider 66-6
Add/Edit CTL Provider 66-7
Configure TLS Proxy Pane 66-8
Adding a TLS Proxy Instance 66-9
Add TLS Proxy Instance Wizard – Server Configuration 66-9
Add TLS Proxy Instance Wizard – Client Configuration 66-10
Add TLS Proxy Instance Wizard – Other Steps 66-12
Edit TLS Proxy Instance – Server Configuration 66-13

Edit TLS Proxy Instance – Client Configuration 66-14
TLS Proxy 66-15
Feature History for the TLS Proxy for Encrypted Voice Inspection 66-17

CHAPTER 67 **Configuring Cisco Mobility Advantage** 67-1
Information about the Cisco Mobility Advantage Proxy Feature 67-1
Cisco Mobility Advantage Proxy Functionality 67-1
Mobility Advantage Proxy Deployment Scenarios 67-2
Trust Relationships for Cisco UMA Deployments 67-4
Licensing for the Cisco Mobility Advantage Proxy Feature 67-6
Configuring Cisco Mobility Advantage 67-6
Task Flow for Configuring Cisco Mobility Advantage 67-7
Feature History for Cisco Mobility Advantage 67-7

CHAPTER 68 **Configuring Cisco Unified Presence** 68-1
Information About Cisco Unified Presence 68-1
Architecture for Cisco Unified Presence for SIP Federation Deployments 68-1
Trust Relationship in the Presence Federation 68-4
Security Certificate Exchange Between Cisco UP and the Security Appliance 68-5
XMPP Federation Deployments 68-5
Configuration Requirements for XMPP Federation 68-6
Licensing for Cisco Unified Presence 68-7
Configuring Cisco Unified Presence Proxy for SIP Federation 68-8
Task Flow for Configuring Cisco Unified Presence Federation Proxy for SIP Federation 68-9
Feature History for Cisco Unified Presence 68-9

CHAPTER 69 **Configuring Cisco Intercompany Media Engine Proxy** 69-1
Information About Cisco Intercompany Media Engine Proxy 69-1
Features of Cisco Intercompany Media Engine Proxy 69-1
How the UC-IME Works with the PSTN and the Internet 69-2
Tickets and Passwords 69-3
Call Fallback to the PSTN 69-5
Architecture and Deployment Scenarios for Cisco Intercompany Media Engine 69-5
Licensing for Cisco Intercompany Media Engine 69-8
Guidelines and Limitations 69-9
Configuring Cisco Intercompany Media Engine Proxy 69-11
Task Flow for Configuring Cisco Intercompany Media Engine 69-11
Configuring NAT for Cisco Intercompany Media Engine Proxy 69-12

Configuring PAT for the Cisco UCM Server 69-14
Creating Access Lists for Cisco Intercompany Media Engine Proxy 69-16
Creating the Media Termination Instance 69-17
Creating the Cisco Intercompany Media Engine Proxy 69-18
Creating Trustpoints and Generating Certificates 69-21
Creating the TLS Proxy 69-24
Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy 69-25
(Optional) Configuring TLS within the Local Enterprise 69-27
(Optional) Configuring Off Path Signaling 69-30
Configuring the Cisco UC-IMC Proxy by using the UC-IME Proxy Pane 69-31
Configuring the Cisco UC-IMC Proxy by using the Unified Communications Wizard 69-33
Feature History for Cisco Intercompany Media Engine Proxy 69-37

PART 12 **Configuring Connection Settings and QoS**

CHAPTER 71 **Configuring Connection Settings** 71-1
Information About Connection Settings 71-1
TCP Intercept and Limiting Embryonic Connections 71-2
Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility 71-2
Dead Connection Detection (DCD) 71-2
TCP Sequence Randomization 71-3
TCP Normalization 71-3
TCP State Bypass 71-3
Licensing Requirements for Connection Settings 71-4
Guidelines and Limitations 71-5
TCP State Bypass Guidelines and Limitations 71-5
Default Settings 71-5
Configuring Connection Settings 71-5
Task Flow For Configuring Configuration Settings (Except Global Timeouts) 71-6
Customizing the TCP Normalizer with a TCP Map 71-6
Configuring Connection Settings 71-8
Configuring Global Timeouts 71-9
Feature History for Connection Settings 71-11

CHAPTER 72 **Configuring QoS** 72-1
Information About QoS 72-1
Supported QoS Features 72-2
What is a Token Bucket? 72-2

Information About Policing 72-3
Information About Priority Queuing 72-3
Information About Traffic Shaping 72-4
How QoS Features Interact 72-4
DSCP and DiffServ Preservation 72-5
Licensing Requirements for QoS 72-5
Guidelines and Limitations 72-5
Configuring QoS 72-6
Determining the Queue and TX Ring Limits for a Standard Priority Queue 72-7
Configuring the Standard Priority Queue for an Interface 72-8
Configuring a Service Rule for Standard Priority Queuing and Policing 72-9
Configuring a Service Rule for Traffic Shaping and Hierarchical Priority Queuing 72-10
Monitoring QoS 72-11
Viewing QoS Police Statistics 72-12
Viewing QoS Standard Priority Statistics 72-12
Viewing QoS Shaping Statistics 72-12
Viewing QoS Standard Priority Queue Statistics 72-13
Feature History for QoS 72-14

CHAPTER 73 **Troubleshooting Connections and Resources** 73-1
Testing Your Configuration 73-1
Pinging ASA Interfaces 73-1
Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping 73-3
Determining Packet Routing with Traceroute 73-6
Tracing Packets with Packet Tracer 73-7
Monitoring Performance 73-8
Monitoring System Resources 73-9
Blocks 73-9
CPU 73-10
Memory 73-10
Monitoring Connections 73-11
Monitoring Per-Process CPU Usage 73-12

PART 13 **Configuring Advanced Network Protection**

CHAPTER 75 **Configuring the ASA for Cisco Cloud Web Security** 75-1
Information About Cisco Cloud Web Security 75-2
Redirection of Web Traffic to Cloud Web Security 75-2

User Authentication and Cloud Web Security 75-2
Authentication Keys 75-3
ScanCenter Policy 75-4
Cloud Web Security Actions 75-5
Bypassing Scanning with Whitelists 75-5
IPv4 and IPv6 Support 75-6
Failover from Primary to Backup Proxy Server 75-6
Licensing Requirements for Cisco Cloud Web Security 75-7
Prerequisites for Cloud Web Security 75-7
Guidelines and Limitations 75-7
Default Settings 75-8
Configuring Cisco Cloud Web Security 75-8
Configuring Communication with the Cloud Web Security Proxy Server 75-8
(Multiple Context Mode) Allowing Cloud Web Security Per Security Context 75-9
Configuring a Service Policy to Send Traffic to Cloud Web Security 75-10
(Optional) Configuring Whitelisted Traffic 75-23
(Optional) Configuring the User Identity Monitor 75-25
Configuring the Cloud Web Security Policy 75-26
Monitoring Cloud Web Security 75-26
Related Documents 75-27
Feature History for Cisco Cloud Web Security 75-27

CHAPTER 76 **Configuring the Botnet Traffic Filter** 76-1
Information About the Botnet Traffic Filter 76-1
Botnet Traffic Filter Address Types 76-2
Botnet Traffic Filter Actions for Known Addresses 76-2
Botnet Traffic Filter Databases 76-2
How the Botnet Traffic Filter Works 76-5
Licensing Requirements for the Botnet Traffic Filter 76-6
Guidelines and Limitations 76-6
Default Settings 76-6
Configuring the Botnet Traffic Filter 76-6
Task Flow for Configuring the Botnet Traffic Filter 76-7
Configuring the Dynamic Database 76-7
Adding Entries to the Static Database 76-9
Enabling DNS Snooping 76-9
Enabling Traffic Classification and Actions for the Botnet Traffic Filter 76-10
Blocking Botnet Traffic Manually 76-12

Searching the Dynamic Database 76-13
Monitoring the Botnet Traffic Filter 76-14
Botnet Traffic Filter Syslog Messaging 76-14
Botnet Traffic Filter Monitor Panes 76-15
Where to Go Next 76-16
Feature History for the Botnet Traffic Filter 76-16

CHAPTER 77 **Configuring Threat Detection** 77-1
Information About Threat Detection 77-1
Licensing Requirements for Threat Detection 77-1
Configuring Basic Threat Detection Statistics 77-2
Information About Basic Threat Detection Statistics 77-2
Guidelines and Limitations 77-3
Default Settings 77-3
Configuring Basic Threat Detection Statistics 77-4
Monitoring Basic Threat Detection Statistics 77-4
Feature History for Basic Threat Detection Statistics 77-5
Configuring Advanced Threat Detection Statistics 77-5
Information About Advanced Threat Detection Statistics 77-5
Guidelines and Limitations 77-5
Default Settings 77-6
Configuring Advanced Threat Detection Statistics 77-6
Monitoring Advanced Threat Detection Statistics 77-7
Feature History for Advanced Threat Detection Statistics 77-8
Configuring Scanning Threat Detection 77-8
Information About Scanning Threat Detection 77-9
Guidelines and Limitations 77-9
Default Settings 77-10
Configuring Scanning Threat Detection 77-10
Feature History for Scanning Threat Detection 77-11

CHAPTER 78 **Using Protection Tools** 78-1
Preventing IP Spoofing 78-1
Configuring the Fragment Size 78-2
Show Fragment 78-2
Configuring TCP Options 78-3
TCP Reset Settings 78-4
Configuring IP Audit for Basic IPS Support 78-5

IP Audit Policy 78-5
Add/Edit IP Audit Policy Configuration 78-5
IP Audit Signatures 78-6
IP Audit Signature List 78-6

CHAPTER 79 **Configuring Filtering Services** 79-1
Information About Web Traffic Filtering 79-1
Filtering URLs and FTP Requests with an External Server 79-2
Information About URL Filtering 79-2
Licensing Requirements for URL Filtering 79-3
Guidelines and Limitations for URL Filtering 79-3
Identifying the Filtering Server 79-3
Configuring Additional URL Filtering Settings 79-4
Configuring Filtering Rules 79-6
Filtering the Rule Table 79-11
Defining Queries 79-12
Feature History for URL Filtering 79-12

PART 14 **Configuring Modules**

CHAPTER 81 **Configuring the ASA IPS Module** 81-1
Information About the ASA IPS module 81-1
How the ASA IPS module Works with the ASA 81-2
Operating Modes 81-3
Using Virtual Sensors (ASA 5510 and Higher) 81-3
Information About Management Access 81-4
Licensing Requirements for the ASA IPS module 81-5
Guidelines and Limitations 81-5
Default Settings 81-6
Configuring the ASA IPS module 81-7
Task Flow for the ASA IPS Module 81-7
Connecting the ASA IPS Management Interface 81-8
Sessioning to the Module from the ASA (May Be Required) 81-11
(ASA 5512-X through ASA 5555-X) Booting the Software Module 81-12
Configuring Basic IPS Module Network Settings 81-12
Configuring the Security Policy on the ASA IPS Module 81-15
Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher) 81-17
Diverting Traffic to the ASA IPS module 81-18

Managing the ASA IPS module 81-19
Monitoring the ASA IPS module 81-24
Feature History for the ASA IPS module 81-24

CHAPTER 82 **Configuring the ASA CX Module** 82-1
Information About the ASA CX Module 82-1
How the ASA CX Module Works with the ASA 82-2
Information About ASA CX Management 82-2
Information About Authentication Proxy 82-3
Information About VPN and the ASA CX Module 82-4
Compatibility with ASA Features 82-4
Licensing Requirements for the ASA CX Module 82-4
Guidelines and Limitations 82-4
Default Settings 82-5
Configuring the ASA CX Module 82-5
Task Flow for the ASA CX Module 82-6
Connecting the ASA CX Management Interface 82-7
Configuring the ASA CX Management IP Address 82-8
Configuring Basic ASA CX Settings at the ASA CX CLI 82-9
Configuring the Security Policy on the ASA CX Module Using PRSM 82-11
(Optional) Configuring the Authentication Proxy Port 82-12
Redirecting Traffic to the ASA CX Module 82-12
Managing the ASA CX Module 82-14
Resetting the Password 82-14
Reloading or Resetting the Module 82-15
Shutting Down the Module 82-15
Monitoring the ASA CX Module 82-15
Showing Module Status 82-16
Showing Module Statistics 82-16
Monitoring Module Connections 82-17
Capturing Module Traffic 82-20
Troubleshooting the ASA CX Module 82-20
Problems with the Authentication Proxy 82-20
Feature History for the ASA CX Module 82-21

CHAPTER 83 **Configuring the ASA CSC Module** 83-1
Information About the CSC SSM 83-1
Determining What Traffic to Scan 83-3

Licensing Requirements for the CSC SSM 83-5
Prerequisites for the CSC SSM 83-5
Guidelines and Limitations 83-6
Default Settings 83-6
Configuring the CSC SSM 83-7
Before Configuring the CSC SSM 83-7
Connecting to the CSC SSM 83-8
Determining Service Policy Rule Actions for CSC Scanning 83-9
CSC SSM Setup Wizard 83-10
Activation/License 83-11
IP Configuration 83-11
Host/Notification Settings 83-12
Management Access Host/Networks 83-13
Password 83-13
Restoring the Default Password 83-14
Wizard Setup 83-15
Using the CSC SSM GUI 83-20
Web 83-20
Mail 83-21
SMTP Tab 83-21
POP3 Tab 83-22
File Transfer 83-22
Updates 83-23
Monitoring the CSC SSM 83-24
Threats 83-24
Live Security Events 83-25
Live Security Events Log 83-25
Software Updates 83-26
Resource Graphs 83-27
Troubleshooting the CSC Module 83-27
Additional References 83-31
Feature History for the CSC SSM 83-31

PART 15 **Configuring VPN**

CHAPTER 85 **VPN Wizards** 85-1
VPN Overview 85-1
IPsec IKEv1 Remote Access Wizard 85-2

Remote Access Client 85-2
VPN Client Authentication Method and Tunnel Group Name 85-3
Client Authentication 85-3
User Accounts 85-4
Address Pool 85-4
Attributes Pushed to Client (Optional) 85-4
IKE Policy 85-5
IPsec Settings (Optional) 85-6
Summary 85-6
IPsec Site-to-Site VPN Wizard 85-7
Peer Device Identification 85-7
IKE Version 85-7
Traffic to Protects 85-7
Authentication Methods 85-7
Encryption Algorithm 85-8
Miscellaneous 85-8
Summary 85-9
AnyConnect VPN Wizard 85-9
Connection Profile Identification 85-9
VPN Protocols 85-9
Client Images 85-10
Authentication Methods 85-10
Client Address Assignment 85-11
Network Name Resolution Servers 85-11
NAT Exempt 85-11
AnyConnect Client Deployment 85-11
Summary 85-12
Clientless SSL VPN Wizard 85-12
SSL VPN Interface 85-12
User Authentication 85-12
Group Policy 85-13
Bookmark List 85-13
Summary 85-13

CHAPTER 86 **Configuring IKE, Load Balancing, and NAC 86-1**
Setting IKE Parameters 86-1
Creating IKE Policies 86-5
Add/Edit IKEv1 Policy 86-6
Add/Edit IKEv2 Policy (Proposal) 86-7

Assignment Policy 86-9
Configuring IPsec 86-10
Adding Crypto Maps 86-10
Pre-Fragmentation 86-17
IPsec Transform Sets 86-18
Configuring Load Balancing 86-20
Creating Virtual Clusters 86-21
Geographical Load Balancing 86-22
Mixed Cluster Scenarios 86-22
Comparing Load Balancing to Failover 86-23
Load Balancing Licensing Requirements 86-24
Eligible Clients 86-24
Load Balancing Prerequisites 86-24
Configuring VPN Cluster Load Balancing with the High Availability and Scalability Wizard 86-25
Configuring Load Balancing (Without the Wizard) 86-26
Enable Clientless SSL VPN Load Balancing Using FQDNs 86-28
Setting Global NAC Parameters 86-29
Configuring Network Admission Control Policies 86-30

CHAPTER 87 **General VPN Setup** 87-1
Client Software 87-1
Edit Client Update Entry 87-3
Default Tunnel Gateway 87-4
Group Policies 87-4
Configuring External Group Policies 87-5
Configuring Network (Client) Access Internal Group Policies 87-7
Configuring Clientless SSL VPN Internal Group Policies 87-23
Configuring Site-to-Site Internal Group Policies 87-28
Defining Time Ranges 87-29
Access Control List Manager 87-31
Standard Access Control List 87-31
Extended Access Control List 87-31
Configuring AnyConnect VPN Client Connections 87-42
Using AnyConnect Client Profiles 87-45
Exempting AnyConnect Traffic from Network Address Translation 87-46
Configuring AnyConnect VPN Connections 87-52
Specifying a Device Certificate for AnyConnect Access 87-53
Configuring Port Settings 87-54

Setting the Basic Attributes for an AnyConnect VPN Connection 87-54
Setting Advanced Attributes for a Connection Profile 87-56
Setting General Attributes for an AnyConnect SSL VPN Connection 87-56
Setting Client Addressing Attributes for an AnyConnect SSL VPN Connection 87-58
Configuring Authentication Attributes for a Connection Profile 87-58
Configuring Secondary Authentication Attributes for an SSL VPN Connection Profile 87-59
Configuring Authorization Attributes for an SSL VPN Connection Profile 87-61
Adding or Editing Content to a Script for Certificate Pre-Fill-Username 87-62
Configuring AnyConnect Secure Mobility 87-64
Add or Edit MUS Access Control 87-66
Configuring Clientless SSL VPN Connections 87-66
Add or Edit Clientless SSL VPN Connections 87-67
Add or Edit Clientless SSL VPN Connections > Basic 87-67
Add or Edit Clientless SSL VPN Connections > Advanced 87-68
Add or Edit Clientless SSL VPN Connections > Advanced > General 87-68
Add or Edit Clientless or SSL VPN Client Connection Profile or IPsec Connection Profiles> Advanced > Authentication 87-69
Assign Authentication Server Group to Interface 87-69
Add or Edit SSL VPN Connections > Advanced > Authorization 87-69
Assign Authorization Server Group to Interface 87-70
Add or Edit SSL VPN Connections > Advanced > SSL VPN 87-70
Add or Edit Clientless SSL VPN Connections > Advanced > Clientless SSL VPN 87-71
Add or Edit Clientless SSL VPN Connections > Advanced > NetBIOS Servers 87-72
Configure DNS Server Groups 87-73
Add or Edit Clientless SSL VPN Connections > Advanced > Clientless SSL VPN 87-73
IPsec Remote Access Connection Profiles 87-73
Add or Edit an IPsec Remote Access Connection Profile 87-74
Add or Edit IPsec Remote Access Connection Profile Basic 87-74
Mapping Certificates to IPsec or SSL VPN Connection Profiles 87-75
Site-to-Site Connection Profiles 87-79
Add/Edit Site-to-Site Connection 87-80
Adding or Editing a Site-to-Site Tunnel Group 87-81
Crypto Map Entry 87-83
Crypto Map Entry for Static Peer Address 87-84
Managing CA Certificates 87-85
Install Certificate 87-85
Configure Options for CA Certificate 87-85
Add/Edit Remote Access Connections > Advanced > General 87-86
Configuring Client Addressing 87-87

Configuring Internal Group Policy IPsec Client Attributes 87-100
Configuring Client Addressing for SSL VPN Connections 87-101
Assign Address Pools to Interface 87-101
Select Address Pools 87-102
Add or Edit an IP Address Pool 87-102
Authenticating SSL VPN Connections 87-102
System Options 87-102
Zone Labs Integrity Server 87-103
Easy VPN Remote 87-105
Advanced Easy VPN Properties 87-107
AnyConnect Essentials 87-108
DTLS Settings 87-109
SSL VPN Client Settings 87-109
Add/Replace SSL VPN Client Image 87-111
Upload Image 87-111
Add/Edit SSL VPN Client Profiles 87-111
Upload Package 87-112
Bypass Interface Access List 87-112
Configuring AnyConnect Host Scan 87-113
Host Scan Dependencies and System Requirements 87-113
Host Scan Packaging 87-114
Installing and Enabling Host Scan on the ASA 87-115
Other Important Documentation Addressing Host Scan 87-119
Configuring Maximum VPN Sessions 87-119
Configuring the Pool of Cryptographic Cores 87-119

CHAPTER 88 **Configuring IP Addresses for VPNs** 88-1

Configuring an IP Address Assignment Policy 88-1
Configuring IPv4 and IPv6 Address Assignments using ASDM 88-2
Viewing Address Assignment Methods 88-3
Configuring Local IP Address Pools 88-3
Configuring Local IPv4 Address Pools Using ASDM 88-3
Configuring Local IPv6 Address Pools Using ASDM 88-4
Configuring AAA Addressing 88-5
Configuring DHCP Addressing 88-6
Configuring DHCP Addressing Using ASDM 88-6

CHAPTER 89 **Configuring Dynamic Access Policies** 89-1

Information About Dynamic Access Policies 89-1
DAP and Endpoint Security 89-2
DAP Support for Remote Access Connection Types 89-2
Remote Access Connection Sequence with DAPs 89-2
Licensing Requirements for Dynamic Access Policies 89-3
Advanced Endpoint Assessment license 89-3
SSL VPN license (client) 89-3
AnyConnect Mobile License 89-3
Dynamic Access Policies Interface 89-4
Configuring Dynamic Access Policies 89-6
Testing Dynamic Access Policies 89-8
DAP and Authentication, Authorization, and Accounting Services 89-9
Configuring AAA Attributes in a DAP 89-9
Retrieving Active Directory Groups 89-11
Configuring Endpoint Attributes Used in DAPs 89-13
Adding an Anti-Spyware or Anti-Virus Endpoint Attribute to a DAP 89-14
Adding an Application Attribute to a DAP 89-15
Adding Mobile Posture Attributes to a DAP 89-16
Adding a File Endpoint Attribute to a DAP 89-17
Adding a Device Endpoint Attribute to a DAP 89-18
Adding a NAC Endpoint Attribute to a DAP 89-19
Adding an Operating System Endpoint Attribute to a DAP 89-20
Adding a Personal Firewall Endpoint Attribute to a DAP 89-20
Adding a Policy Endpoint Attribute to a DAP 89-21
Adding a Process Endpoint Attribute to a DAP 89-22
Adding a Registry Endpoint Attribute to a DAP 89-23
DAP and AntiVirus, AntiSpyware, and Personal Firewall Programs 89-24
Endpoint Attribute Definitions 89-24
Configuring DAP Access and Authorization Policy Attributes 89-27
Performing a DAP Trace 89-31
Guide to Creating DAP Logical Expressions using LUA 89-31
Syntax for Creating Lua EVAL Expressions 89-32
The DAP CheckAndMsg Function 89-33
Additional Lua Functions 89-35
CheckAndMsg with Custom Function Example 89-38
Further Information on Lua 89-38
Operator for Endpoint Category 89-38

DAP Examples 89-38

CHAPTER 90 **Clientless SSL VPN End User Set-up** 90-1

Requiring Usernames and Passwords 90-1

Communicating Security Tips 90-2

Configuring Remote Systems to Use Clientless SSL VPN Features 90-2

Capturing Clientless SSL VPN Data 90-7

Creating a Capture File 90-7

Using a Browser to Display Capture Data 90-8

CHAPTER 91 **Configuring Clientless SSL VPN** 91-1

Information About Clientless SSL VPN 91-2

Licensing Requirements 91-2

Prerequisites for Clientless SSL VPN 91-5

Guidelines and Limitations 91-5

Observing Clientless SSL VPN Security Precautions 91-5

Configuring Clientless SSL VPN Access 91-6

Clientless SSL VPN Server Certificate Verification 91-7

Configuring ACLs 91-11

Configuring the Setup for Cisco Secure Desktop 91-13

Configuring Application Profile Customization Framework 91-15

Restrictions 91-15

Managing APCF Profiles 91-15

Uploading APCF Packages 91-16

Managing APCF Packets 91-16

APCF Syntax 91-17

Managing Passwords 91-20

Configuring SSO with the HTTP Form Protocol 91-22

Using Auto Signon 91-25

Accessing Virtual Desktop Infrastructure (VDI) 91-27

Configuring Session Settings 91-28

Java Code Signer 91-28

Encoding 91-29

Content Cache 91-30

Content Rewrite 91-31

Configuration Example for Content Rewrite Rules 91-32

Configuring Browser Access to Plug-ins 91-33

Preparing the Security Appliance for a Plug-in 91-34
Installing Plug-ins Redistributed By Cisco 91-34
Providing Access to a Citrix XenApp Server 91-36
Microsoft Kerberos Constrained Delegation Solution 91-38
Understanding How KCD Works 91-38
Configuring Application Access 91-45
Configuring Smart Tunnel Access 91-45
Configuring Smart Tunnel Log Off 91-54
Configuring Port Forwarding 91-55
Information About Port Forwarding 91-55
Configuring DNS for Port Forwarding 91-57
Adding Applications to Be Eligible for Port Forwarding 91-59
Adding/Editing Port Forwarding Entry 91-59
Assigning a Port Forwarding List 91-59
Enabling and Disabling Port Forwarding 91-60
Configuring the Use of External Proxy Servers 91-60
SSO Servers 91-61
Configuring SiteMinder and SAML Browser Post Profile 91-62
Application Access User Notes 91-64
Closing Application Access to Prevent hosts File Errors 91-64
Recovering from hosts File Errors When Using Application Access 91-64
Configuring File Access 91-68
CIFS File Access Requirement and Limitation 91-68
Ensuring Clock Accuracy for SharePoint Access 91-69
Customizing the Clientless SSL VPN User Experience 91-69
Customizing the Logon Page with the Customization Editor 91-69
Replacing the Logon Page with your own Fully Customized Page 91-71
Using Clientless SSL VPN with PDAs 91-74
Using E-Mail over Clientless SSL VPN 91-74
Configuring E-mail Proxies 91-75
Configuring Web E-mail: MS Outlook Web App 91-75
Configuring Portal Access Rules 91-75
Clientless SSL VPN End User Setup 91-77
Defining the End User Interface 91-77
Customizing Clientless SSL VPN Pages 91-80
Information About Customization 91-80
Exporting a Customization Template 91-81
Editing the Customization Template 91-81

Customizing the Portal Page 91-90
Customizing the Logout Page 91-92
Customizing the External Portal Page 91-93
Adding Customization Object 91-93
Importing/Exporting Customization Object 91-94
Creating XML-Based Portal Customization Objects and URL Lists 91-94
Understanding the XML Customization File Structure 91-95
Configuration Example for Customization 91-98
Using the Customization Template 91-101
Help Customization 91-113
Import/Export Application Help Content 91-116
Configuring Browser Access to Client-Server Plug-ins 91-118
About Installing Browser Plug-ins 91-118
Preparing the Security Appliance for a Plug-in 91-120
Customizing Help 91-121
Communicating Security Tips 91-123
Configuring Remote Systems to Use Clientless SSL VPN Features 91-123
Starting Clientless SSL VPN 91-124
Using the Clientless SSL VPN Floating Toolbar 91-124
Browsing the Web 91-125
Browsing the Network (File Management) 91-125
Using Port Forwarding 91-127
Using E-mail Via Port Forwarding 91-128
Using E-mail Via Web Access 91-129
Using E-mail Via E-mail Proxy 91-129
Using Smart Tunnel 91-129
Customizing the AnyConnect Client 91-131
Customizing AnyConnect by Importing Resource Files 91-132
Customizing Your Own AnyConnect GUI Text and Scripts 91-133
Customizing AnyConnect GUI Text and Messages 91-136
Customizing the Installer Program Using Installer Transforms 91-137
Localizing the Install Program using Installer Transforms 91-138
Importing/Exporting Language Localization 91-139
Configuring Bookmarks 91-140
Adding a Bookmark for a URL with a GET or Post Method 91-141
Adding a URL for a Predefined Application Template 91-142
Adding a Bookmark for an Auto Sign-On Application 91-144
Importing/Exporting Bookmark List 91-145
Importing/Exporting GUI Customization Objects (Web Contents) 91-145

Adding/Editing Post Parameter 91-146
Sending an Administrator's Alert to Clientless SSL VPN Users 91-149

CHAPTER 92 E-Mail Proxy 92-1
Configuring E-Mail Proxy 92-1
AAA 92-2
POP3S Tab 92-2
IMAP4S Tab 92-4
SMTPS Tab 92-5
Access 92-7
Edit E-Mail Proxy Access 92-8
Authentication 92-8
Default Servers 92-10
Delimiters 92-11

CHAPTER 93 Monitoring VPN 93-1
VPN Connection Graphs 93-1
IPsec Tunnels 93-1
Sessions 93-2
VPN Statistics 93-2
Sessions Window 93-2
Viewing Active AnyConnect Sessions 93-5
Viewing VPN Sessions Details 93-6
Cluster Loads 93-8
Crypto Statistics 93-9
Compression Statistics 93-9
Encryption Statistics 93-9
Global IKE/IPsec Statistics 93-10
NAC Session Summary 93-10
Protocol Statistics 93-11
VLAN Mapping Sessions 93-11
SSO Statistics for Clientless SSL VPN Session 93-11
VPN Connection Status for the Easy VPN Client 93-13

CHAPTER 94 Configuring SSL Settings 94-1
SSL 94-1
Edit SSL Certificate 94-2
SSL Certificates 94-2

PART 16 Configuring Logging, SNMP, and Smart Call Home

CHAPTER 96 Configuring Logging 96-1

Information About Logging 96-1
Logging in Multiple Context Mode 96-2
Analyzing Syslog Messages 96-2
Syslog Message Format 96-3
Severity Levels 96-3
Message Classes and Range of Syslog IDs 96-4
Filtering Syslog Messages 96-4
Sorting in the Log Viewers 96-4
Using Custom Message Lists 96-5
Using Clustering 96-5
Licensing Requirements for Logging 96-5
Prerequisites for Logging 96-6
Guidelines and Limitations 96-6
Configuring Logging 96-6
Enabling Logging 96-7
Configuring an Output Destination 96-7
Monitoring the Logs 96-24
Filtering Syslog Messages Through the Log Viewers 96-25
Editing Filtering Settings 96-27
Executing Certain Commands Using the Log Viewers 96-27
Feature History for Logging 96-28

CHAPTER 97 Configuring NetFlow Secure Event Logging (NSEL) 97-1

Information About NSEL 97-1
Using NSEL and Syslog Messages 97-2
Using NSEL in Clustering 97-3
Licensing Requirements for NSEL 97-4
Prerequisites for NSEL 97-4
Guidelines and Limitations 97-4
Configuring NSEL 97-5
Using NetFlow 97-5
Matching NetFlow Events to Configured Collectors 97-6
Monitoring NSEL 97-7
Where to Go Next 97-7
Additional References 97-7

Related Documents 97-8
RFCs 97-8
Feature History for NSEL 97-8

CHAPTER 98 **Configuring SNMP** 98-1
Information About SNMP 98-1
Information About SNMP Terminology 98-2
SNMP Version 3 98-2
Licensing Requirements for SNMP 98-4
Prerequisites for SNMP 98-4
Guidelines and Limitations 98-4
Configuring SNMP 98-6
Enabling SNMP 98-6
Configuring an SNMP Management Station 98-6
Configuring SNMP Traps 98-7
Using SNMP Version 1 or 2c 98-8
Using SNMP Version 3 98-8
Monitoring SNMP 98-9
SNMP Syslog Messaging 98-10
SNMP Monitoring 98-10
Where to Go Next 98-11
Additional References 98-11
RFCs for SNMP Version 3 98-11
MIBs 98-11
Application Services and Third-Party Tools 98-13
Feature History for SNMP 98-13

CHAPTER 99 **Configuring Anonymous Reporting and Smart Call Home** 99-1
Information About Anonymous Reporting and Smart Call Home 99-1
Information About Anonymous Reporting 99-2
Information About Smart Call Home 99-4
Licensing Requirements for Anonymous Reporting and Smart Call Home 99-4
Prerequisites for Smart Call Home and Anonymous Reporting 99-5
Guidelines and Limitations 99-5
Configuring Anonymous Reporting and Smart Call Home 99-6
Configuring Anonymous Reporting 99-6
Configuring Smart Call Home 99-6
Monitoring Anonymous Reporting and Smart Call Home 99-10

Feature History for Anonymous Reporting and Smart Call Home 99-11

PART 17 System Administration

CHAPTER 101 Managing Software and Configurations 101-1

Upgrading the Software 101-1
Viewing Your Current Version 101-2
Downloading the Software from Cisco.com 101-2
Upgrading a Standalone Unit 101-2
Upgrading a Failover Pair or ASA Cluster 101-7
Managing Files 101-14
Accessing the File Management Tool 101-14
Managing Mount Points 101-15
Transferring Files 101-17
Configuring the Images and Startup Configuration to Use 101-19
Backing Up and Restoring Configurations or Other Files 101-20
Backing Up Configurations 101-21
Backing Up the Local CA Server 101-24
Restoring Configurations 101-25
Saving the Running Configuration to a TFTP Server 101-27
Scheduling a System Restart 101-27
Downgrading Your Software 101-28
Information About Activation Key Compatibility 101-28
Performing the Downgrade 101-29
Configuring Auto Update 101-30
Information About Auto Update 101-30
Configuring Communication with an Auto Update Server 101-31

CHAPTER 102 Troubleshooting 102-1

Configuring and Running Captures with the Packet Capture Wizard 102-1
Ingress Traffic Selector 102-3
Egress Traffic Selector 102-4
Buffers 102-4
Summary 102-4
Run Captures 102-4
Save Captures 102-5

PART 18 Reference

APPENDIX A **Addresses, Protocols, and Ports** A-1

IPv4 Addresses and Subnet Masks A-1

Classes A-1

Private Networks A-2

Subnet Masks A-2

IPv6 Addresses A-5

IPv6 Address Format A-5

IPv6 Address Types A-6

IPv6 Address Prefixes A-10

Protocols and Applications A-11

TCP and UDP Ports A-11

Local Ports and Protocols A-14

ICMP Types A-15

APPENDIX B **Configuring an External Server for Authorization and Authentication** B-1

Understanding Policy Enforcement of Permissions and Attributes B-1

Configuring an External LDAP Server B-2

Organizing the ASA for LDAP Operations B-3

Defining the ASA LDAP Configuration B-5

Active Directory/LDAP VPN Remote Access Authorization Examples B-15

Configuring an External RADIUS Server B-25

Reviewing the RADIUS Configuration Procedure B-26

ASA RADIUS Authorization Attributes B-26

ASA IETF RADIUS Authorization Attributes B-36

RADIUS Accounting Disconnect Reason Codes B-36

Configuring an External TACACS+ Server B-37

GLOSSARY

INDEX

About This Guide

This preface introduces *Cisco ASA Series ASDM Configuration Guide* and includes the following sections:

- Document Objectives, page lxi
- Audience, page lxi
- Related Documentation, page lxii
- Conventions, page lxii
- Obtaining Documentation and Submitting a Service Request, page lxiii

Document Objectives

The purpose of this guide is to help you configure the ASA using ASDM. This guide does not cover every feature, but describes only the most common configuration scenarios.

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

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

Audience

This guide is for network managers who perform any of the following tasks:

- Manage network security
- Install and configure firewalls/ASAs
- Configure VPNs
- Configure intrusion detection software

Related Documentation

For more information, see *Navigating the Cisco ASA Series Documentation* at http://www.cisco.com/en/US/docs/security/asa/roadmap/asaroadmap.html.

Conventions

This document uses the following conventions:

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

Note Means *reader take note*.

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

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

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html.

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

PART 1

Getting Started with the ASA

Introduction to the Cisco ASA

The Cisco ASA provides advanced stateful firewall and VPN concentrator functionality in one device, and for some models, integrated services modules such as IPS. The ASA includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), clustering (combining multiple firewalls into a single firewall), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, IPsec VPN, SSL VPN, and clientless SSL VPN support, and many more features.

Note ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Similarly, if a feature was added into a maintenance release for an older major or minor version, then the ASDM documentation includes the new feature even though that feature might not be available in all later ASA releases. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see *Cisco ASA 5500 Series Hardware and Software Compatibility*.

This chapter includes the following sections:

- ASDM Client Operating System and Browser Requirements, page 1-1
- Hardware and Software Compatibility, page 1-4
- VPN Specifications, page 1-4
- New Features, page 1-5
- How the ASA Services Module Works with the Switch, page 1-19
- Firewall Functional Overview, page 1-20
- VPN Functional Overview, page 1-25
- Security Context Overview, page 1-25
- ASA Clustering Overview, page 1-26

ASDM Client Operating System and Browser Requirements

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

Table 1-1 Operating System and Browser Requirements

Operating System	Browser				Java SE Plug-in
	Internet Explorer	Firefox	Safari	Chrome	
Microsoft Windows (English and Japanese): • 7 • Vista • 2008 Server • XP	6.0 or later	1.5 or later	No support	18.0 or later	6.0 or later
Apple Macintosh OS X: • 10.8 • 10.7 • 10.6 • 10.5 • 10.4	No support	1.5 or later	2.0 or later	18.0 or later	6.0 or later
Red Hat Enterprise Linux 5 (GNOME or KDE): • Desktop • Desktop with Workstation	N/A	1.5 or later	N/A	18.0 or later	6.0 or later

See the following caveats:

- If you upgrade from a previous version to Java 7 update 5, you may not be able to open ASDM using the Java Web Start from an IPv6 address; you can either download the ASDM Launcher, or follow the instructions at: http://java.com/en/download/help/clearcache_upgrade.xml.
- ASDM requires you to make an SSL connection to the ASA in the following situations:
  - When you first connect your browser to the ASA and access the ASDM splash screen.
  - When you launch ASDM using the launcher or the Java web start application.

  If the ASA only has the base encryption license (DES), and therefore has weak encryption ciphers for the SSL connection, you may not be able to access the splash screen or launch ASDM. See the following issues:
  - When using Java 7 when launching ASDM, you must have the strong encryption license (3DES/AES) on the ASA. With only the base encryption license (DES), you cannot launch ASDM. Even if you can connect with a browser to the ASDM splash screen and download the launcher or web start application, you cannot then launch ASDM. You must uninstall Java 7, and install Java 6.
  - When using Java 6 for accessing the splash screen in a browser, by default, Internet Explorer on Windows Vista and later and Firefox on all operating systems do not support DES for SSL; therefore without the strong encryption license (3DES/AES), see the following workarounds:

If available, use an already downloaded ASDM launcher or Java web start application. The launcher works with Java 6 and weak encryption, even if the browsers do not.

For Windows Internet Explorer, you can enable DES as a workaround. See http://support.microsoft.com/kb/929708 for details.

For Firefox on any operating system, you can enable the security.ssl3.dhe_dss_des_sha setting as a workaround. See http://kb.mozillazine.org/About:config to learn how to change hidden configuration preferences.

- When the ASA uses a self-signed certificate or an untrusted certificate, Firefox 4 and later and Safari are unable to add security exceptions when browsing using HTTPS over IPv6. See: https://bugzilla.mozilla.org/show_bug.cgi?id=633001. This caveat affects all SSL connections originating from Firefox or Safari to the ASA (including ASDM connections). To avoid this caveat, configure a proper certificate for the ASA that is issued by a trusted certificate authority.
- If you change the SSL encryption on the ASA to exclude both RC4-MD5 and RC4-SHA1 algorithms (these algorithms are enabled by default), then Chrome cannot launch ASDM due to the Chrome "SSL false start" feature. We suggest re-enabling one of these algorithms (see the Configuration > Device Management > Advanced > SSL Settings pane); or you can disable SSL false start in Chrome using the **--disable-ssl-false-start** flag according to http://www.chromium.org/developers/how-tos/run-chromium-with-flags.
- For Internet Explorer 9.0 for servers, the "Do not save encrypted pages to disk" option is enabled by default (See Tools > Internet Options > Advanced). This option causes the initial ASDM download to fail. Be sure to disable this option to allow ASDM to download.
- On MacOS, you may be prompted to install Java the first time you run ASDM; follow the prompts as necessary. ASDM will launch after the installation completes.
- On MacOS, you may see the following error message when opening the ASDM Launcher:

```
Cannot launch Cisco ASDM-IDM. No compatible version of Java 1.5+ is available.
```

In this case, Java 7 is the currently-preferred Java version; you need to set Java 6 as the preferred Java version: Open the **Java Preferences** application (under Applications > Utilities), select the preferred Java version, and drag it up to be the first line in the table.

- On MacOS 10.8 and later, you need to allow applications that are not signed with an Apple Developer ID. If you do not change your security preferences, you see an error screen.

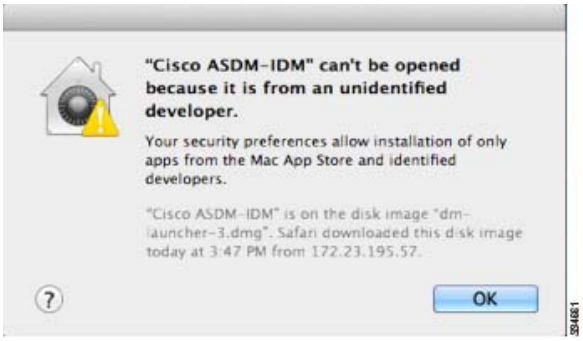

a. To change the security setting, open System Preferences, and click **Security & Privacy**.

b. On the General tab, under Allow applications downloaded from, click **Anywhere**.

Hardware and Software Compatibility

For a complete list of supported hardware and software, see the *Cisco ASA Compatibility*:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

VPN Specifications

See *Supported VPN Platforms, Cisco ASA 5500 Series*:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asa-vpn-compatibility.html

New Features

- New Features in ASA 9.0(2)/ASDM 7.1(2), page 1-5
- New Features in ASA 8.4(5)/ASDM 7.0(2), page 1-5
- New Features in ASA 9.0(1)/ASDM 7.0(1), page 1-5

Note New, changed, and deprecated syslog messages are listed in syslog messages guide.

New Features in ASA 9.0(2)/ASDM 7.1(2)

Released: February 25, 2013

Table 1 lists the new features for ASA Version 9.0(2)/ASDM Version 7.1(2).

Note Features added in 8.4(4.x) and 8.4(5) are not included in 9.0(2) unless they were listed in the 9.0(1) feature table.

Table 1-2 New Features for ASA Version 9.0(2)/ASDM Version 7.1(2)

Feature	Description
Remote Access Features	
Clientless SSL VPN: Windows 8 Support	This release adds support for Windows 8 x86 (32-bit) and Windows 8 x64 (64-bit) operating systems. We support the following browsers on Windows 8: • Internet Explorer 10 (desktop only) • Firefox (all supported Windows 8 versions) • Chrome (all supported Windows 8 versions) See the following limitations: • Internet Explorer 10: – The Modern (AKA Metro) browser is not supported. – If you enable Enhanced Protected Mode, we recommend that you add the ASA to the trusted zone. – If you enable Enhanced Protected Mode, Smart Tunnel and Port Forwarder are not supported. • A Java Remote Desktop Protocol (RDP) plugin connection to a Windows 8 PC is not supported.
Cisco Secure Desktop: Windows 8 Support	CSD 3.6.6215 was updated to enable selection of Windows 8 in the Prelogin Policy operating system check. See the following limitations: • Secure Desktop (Vault) is not supported with Windows 8.
Dynamic Access Policies: Windows 8 Support	ASDM was updated to enable selection of Windows 8 in the DAP Operating System attribute.

New Features in ASA 8.4(5)/ASDM 7.0(2)

Released: October 31, 2012

Table 1 lists the new features for ASA Version 8.4(5)/ASDM Version 7.0(2).

Table 1-3 New Features for ASA Version 8.4(5)/ASDM Version 7.0(2)

Feature	Description
Firewall Features	
EtherType ACL support for IS-IS traffic (transparent firewall mode)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following command: **access-list ethertype** {**permit** | **deny**} **is-is**. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules. *This feature is not available in 8.5(1), 8.6(1), 8.7(1), 9.0(1), or 9.1(1).*
ARP cache additions for non-connected subnets	The ASA ARP cache only contains entries from directly-connected subnets by default. You can now enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries. You may want to use this feature if you use: • Secondary subnets. • Proxy ARP on adjacent routes for traffic forwarding. We modified the following screen: Configuration > Device Management > Advanced > ARP > ARP Static Table. *This feature is not available in 8.5(1), 8.6(1), or 8.7(1).*
Increased maximum connection limits for service policy rules	The maximum number of connections for service policy rules was increased from 65535 to 2000000. We modified the following screen: Configuration > Firewall > Service Policy Rules > Connection Settings. *This feature is not available in 8.5(1), 8.6(1), or 8.7(1).*
Remote Access Features	
Improved Host Scan and ASA Interoperability	Host Scan and the ASA use an improved process to transfer posture attributes from the client to the ASA. This gives the ASA more time to establish a VPN connection with the client and apply a dynamic access policy. *This feature is not available in 8.5(1), 8.6(1), 8.7(1), 9.0(1), or 9.1(1).*
Monitoring Features	
NAT-MIB cnatAddrBindNumberOfEntries and cnatAddrBindSessionCount OIDs to allow polling for Xlate count.	Support was added for the NAT-MIB cnatAddrBindNumberOfEntries and cnatAddrBindSessionCount OIDs to support xlate_count and max_xlate_count for SNMP. This data is equivalent to the **show xlate count** command. *This feature is not available in 8.5(1), 8.6(1), 8.7(1), 9.0(1), or 9.1(1).*

Table 1-3 New Features for ASA Version 8.4(5)/ASDM Version 7.0(2) (continued)

Feature	Description
NSEL	Flow-update events have been introduced to provide periodic byte counters for flow traffic. You can change the time interval at which flow-update events are sent to the NetFlow collector. You can filter to which collectors flow-update records will be sent. We modified the following screens: Configuration > Device Management > Logging > NetFlow. Configuration > Firewall > Service Policy Rules > Add Service Policy Rule Wizard - Rule Actions > NetFlow > Add Flow Event *This feature is not available in 8.5(1), 8.6(1), 8.7(1), 9.0(1), or 9.1(1).*
Hardware Features	
ASA 5585-X DC power supply support	Support was added for the ASA 5585-X DC power supply. *This feature is not available in 8.5(1), 8.6(1), 8.7(1), 9.0(1), or 9.1(1).*

New Features in ASA 9.0(1)/ASDM 7.0(1)

Released: October 29, 2012

Table 1-2 lists the new features for ASA Version 9.0(1)/ASDM Version 7.0(1).

Note Features added in 8.4(4.x) and 8.4(5) are not included in 9.0(1) unless they are explicitly listed in this table.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1)

Feature	Description
Firewall Features	
Cisco TrustSec integration	Cisco TrustSec provides an access-control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec solution, enforcement devices utilize a combination of user attributes and end-point attributes to make role-based and identity-based access control decisions. In this release, the ASA integrates with Cisco TrustSec to provide security group based policy enforcement. Access policies within the Cisco TrustSec domain are topology-independent, based on the roles of source and destination devices rather than on network IP addresses. The ASA can utilize the Cisco TrustSec solution for other types of security group based policies, such as application inspection; for example, you can configure a class map containing an access policy based on a security group. We introduced the following MIB: CISCO-TRUSTSEC-SXP-MIB. We introduced or modified the following screens: Configuration > Firewall > Identity by TrustSec Configuration > Firewall > Objects > Security Groups Object Groups Configuration > Firewall > Access Rules > Add Access Rules Monitoring > Properties > Identity by TrustSec > PAC Monitoring > Properties > Identity by TrustSec > Environment Data Monitoring > Properties > Identity by TrustSec > SXP Connections Monitoring > Properties > Identity by TrustSec > IP Mappings Monitoring > Properties > Connections Tools > Packet Tracer
Cisco Cloud Web Security (ScanSafe)	Cisco Cloud Web Security provides content scanning and other malware protection service for web traffic. It can also redirect and report about web traffic based on user identity. **Note** Clientless SSL VPN is not supported with Cloud Web Security; be sure to exempt any clientless SSL VPN traffic from the ASA service policy for Cloud Web Security. We introduced or modified the following screens: Configuration > Device Management > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit > Manage Cloud Web Security Class Maps Configuration > Firewall > Identity Options Configuration > Firewall > Service Policy Rules Monitoring > Properties > Cloud Web Security

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule
Unified communications support on the ASASM	The ASASM now supports all Unified Communications features.
Per-session PAT	The per-session PAT feature improves the scalability of PAT and, for ASA clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.
ARP cache additions for non-connected subnets	The ASA ARP cache only contains entries from directly-connected subnets by default. You can now enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries. You may want to use this feature if you use: • Secondary subnets. • Proxy ARP on adjacent routes for traffic forwarding. We modified the following screen: Configuration > Device Management > Advanced > ARP > ARP Static Table. *Also available in 8.4(5).*

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
SunRPC change from dynamic ACL to pin-hole mechanism	Previously, Sun RPC inspection does not support outbound access lists because the inspection engine uses dynamic access lists instead of secondary connections. In this release, when you configure dynamic access lists on the ASA, they are supported on the ingress direction only and the ASA drops egress traffic destined to dynamic ports. Therefore, Sun RPC inspection implements a pinhole mechanism to support egress traffic. Sun RPC inspection uses this pinhole mechanism to support outbound dynamic access lists. *Also available in 8.4(4.1).*
Inspection reset action change	Previously, when the ASA dropped a packet due to an inspection engine rule, the ASA sent only one RST to the source device of the dropped packet. This behavior could cause resource issues. In this release, when you configure an inspection engine to use a reset action and a packet triggers a reset, the ASA sends a TCP reset under the following conditions: • The ASA sends a TCP reset to the inside host when the **service resetoutbound** command is enabled. (The **service resetoutbound** command is disabled by default.) • The ASA sends a TCP reset to the outside host when the **service resetinbound** command is enabled. (The **service resetinbound** command is disabled by default.) For more information, see the **service** command in the ASA command reference. This behavior ensures that a reset action will reset the connections on the ASA and on inside servers; therefore countering denial of service attacks. For outside hosts, the ASA does not send a reset by default and information is not revealed through a TCP reset. *Also available in 8.4(4.1).*
Increased maximum connection limits for service policy rules	The maximum number of connections for service policy rules was increased from 65535 to 2000000. We modified the following screen: Configuration > Firewall > Service Policy Rules > Connection Settings. *Also available in 8.4(5)*
High Availability and Scalability Features	

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
ASA Clustering for the ASA 5580 and 5585-X	ASA Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. ASA clustering is supported for the ASA 5580 and the ASA 5585-X; all units in a cluster must be the same model with the same hardware specifications. See the configuration guide for a list of unsupported features when clustering is enabled. We introduced or modified the following screens: Home > Device Dashboard Home > Cluster Dashboard Home > Cluster Firewall Dashboard Configuration > Device Management > Advanced > Address Pools > MAC Address Pools Configuration > Device Management > High Availability and Scalability > ASA Cluster Configuration > Device Management > Logging > Syslog Setup > Advanced Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced Configuration > Device Setup > Interfaces > Add/Edit Interface > IPv6 Configuration > Device Setup > Interfaces > Add/Edit EtherChannel Interface > Advanced Configuration > Firewall > Advanced > Per-Session NAT Rules Monitoring > ASA Cluster Monitoring > Properties > System Resources Graphs > Cluster Control Link Tools > Preferences > General Tools > System Reload Tools > Upgrade Software from Local Computer Wizards > High Availability and Scalability Wizard Wizards > Packet Capture Wizard Wizards > Startup Wizard
OSPF, EIGRP, and Multicast for clustering	For OSPFv2 and OSPFv3, bulk synchronization, route synchronization, and spanned EtherChannels are supported in the clustering environment. For EIGRP, bulk synchronization, route synchronization, and spanned EtherChannels are supported in the clustering environment. Multicast routing supports clustering.
Packet capture for clustering	To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the master unit using the **cluster exec capture** command, which is then automatically enabled on all of the slave units in the cluster. The **cluster exec** keywords are the new keywords that you place in front of the **capture** command to enable cluster-wide capture. We modified the following screen: Wizards > Packet Capture Wizard.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Logging for clustering	Each unit in the cluster generates syslog messages independently. You can use the **logging device-id** command to generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster. We modified the following screen: Configuration > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration.
Configure the connection replication rate during a bulk sync	You can now configure the rate at which the ASA replicates connections to the standby unit when using Stateful Failover. By default, connections are replicated to the standby unit during a 15 second period. However, when a bulk sync occurs (for example, when you first enable failover), 15 seconds may not be long enough to sync large numbers of connections due to a limit on the maximum connections per second. For example, the maximum connections on the ASA is 8 million; replicating 8 million connections in 15 seconds means creating 533 K connections per second. However, the maximum connections allowed per second is 300 K. You can now specify the rate of replication to be less than or equal to the maximum connections per second, and the sync period will be adjusted until all the connections are synchronized. *Also available in 8.4(4.1) and 8.5(1.7).*
IPv6 Features	
IPv6 Support on the ASA's outside interface for VPN Features.	This release of the ASA adds support for IPv6 VPN connections to its outside interface using SSL and IKEv2/IPsec protocols. This release of the ASA continues to support IPv6 VPN traffic on its inside interface using the SSL protocol as it has in the past. This release does not provide IKEv2/IPsec protocol on the inside interface.
Remote Access VPN support for IPv6: IPv6 Address Assignment Policy	You can configure the ASA to assign an IPv4 address, an IPv6 address, or both an IPv4 and an IPv6 address to an AnyConnect client by creating internal pools of addresses on the ASA or by assigning a dedicated address to a local user on the ASA. The endpoint must have the dual-stack protocol implemented in its operating system to be assigned both types of addresses. Assigning an IPv6 address to the client is supported for the SSL protocol. This feature is not supported for the IKEv2/IPsec protocol. We modified the following screens: Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Assignment Policy Configuration > Remote Access VPN > AAA/Local Users > Local Users > (Edit local user account) > VPN Policy

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Remote Access VPN support for IPv6: Assigning DNS Servers with IPv6 Addresses to group policies	DNS servers can be defined in a Network (Client) Access internal group policy on the ASA. You can specify up to four DNS server addresses including up to two IPv4 addresses and up to two IPv6 addresses. DNS servers with IPv6 addresses can be reached by VPN clients when they are configured to use the SSL protocol. This feature is not supported for clients configured to use the IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Network (Client) Access > Group Policies > (Edit group policy) > Servers.
Remote Access VPN support for IPv6: Split tunneling	Split tunneling enables you to route some network traffic through the VPN tunnel (encrypted) and to route other network traffic outside the VPN tunnel (unencrypted or "in the clear"). You can now perform split tunneling on IPv6 network traffic by defining an IPv6 policy which specifies a unified access control rule. IPv6 split tunneling is reported with the telemetric data sent by the Smart Call Home feature. If either IPv4 or IPv6 split tunneling is enabled, Smart Call Home reports split tunneling as "enabled." For telemetric data, the VPN session database displays the IPv6 data typically reported with session management. You can include or exclude IPv6 traffic from the VPN "tunnel" for VPN clients configured to use the SSL protocol. This feature is not supported for the IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Network (Client) Access > Group Policies > (Edit group policy) > Advanced > Split Tunneling.
Remote Access VPN support for IPv6: AnyConnect Client Firewall Rules	Access control rules for client firewalls support access list entries for both IPv4 and IPv6 addresses. ACLs containing IPv6 addresses can be applied to clients configured to use the SSL protocol. This feature is not supported for the IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Network (Client) Access > Group Policies > (Edit group policy) > Advanced > AnyConnect Client > Client Firewall.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Remote Access VPN support for IPv6: Client Protocol Bypass	The Client Protocol Bypass feature allows you to configure how the ASA manages IPv4 traffic when it is expecting only IPv6 traffic or how it manages IPv6 traffic when it is expecting only IPv4 traffic. When the AnyConnect client makes a VPN connection to the ASA, the ASA could assign it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the ASA assigns the AnyConnect connection only an IPv4 address or only an IPv6 address, you can now configure the Client Bypass Protocol to drop network traffic for which the ASA did not assign an IP address, or allow that traffic to bypass the ASA and be sent from the client unencrypted or "in the clear." For example, assume that the ASA assigns only an IPv4 address to an AnyConnect connection and the endpoint is dual stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear. This feature can be used by clients configured to use the SSL or IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Network (Client) Access > Group Policies > (Group Policy) Advanced > AnyConnect Client > Client Bypass Protocol.
Remote Access VPN support for IPv6: IPv6 Interface ID and prefix	You can now specify a dedicated IPv6 address for local VPN users. This feature benefits users configured to use the SSL protocol. This feature is not supported for the IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > AAA/Local Users > Local Users > (Edit User) > VPN Policy.
Remote Access VPN support for IPv6: Sending ASA FQDN to AnyConnect client	You can return the FQDN of the ASA to the AnyConnect client to facilitate load balancing and session roaming. This feature can be used by clients configured to use the SSL or IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Network (Client) Access > Group Policies > (Edit group policy) > Advanced > AnyConnect.
Remote Access VPN support for IPv6: ASA VPN Load Balancing	Clients with IPv6 addresses can make AnyConnect connections through the public-facing IPv6 address of the ASA cluster or through a GSS server. Likewise, clients with IPv6 addresses can make AnyConnect VPN connections through the public-facing IPv4 address of the ASA cluster or through a GSS server. Either type of connection can be load-balanced within the ASA cluster. For clients with IPv6 addresses to successfully connect to the ASAs public-facing IPv4 address, a device that can perform network address translation from IPv6 to IPv4 needs to be in the network. This feature can be used by clients configured to use the SSL or IKEv2/IPsec protocol. We modified the following screen: Configuration > Remote Access VPN > Load Balancing.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Remote Access VPN support for IPv6: Dynamic Access Policies support IPv6 attributes	When using ASA 9.0 or later with ASDM 6.8 or later, you can now specify these attributes as part of a dynamic access policy (DAP): • IPv6 addresses as a Cisco AAA attribute • IPv6 TCP and UDP ports as part of a Device endpoint attribute • Network ACL Filters (client) This feature can be used by clients configured to use the SSL or IKEv2/IPsec protocol. We modified the following screens: Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies > Add > Cisco AAA attribute Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies > Add > Device > Add Endpoint Attribute Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies > Network ACL Filters (client) Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies > Webtype ACL Filters (clientless)
Remote Access VPN support for IPv6: Session Management	Session management output displays the IPv6 addresses in Public/Assigned address fields for AnyConnect connections, site-to-site VPN connections, and Clientless SSL VPN connections. You can add new filter keywords to support filtering the output to show only IPv6 (outside or inside) connections. No changes to IPv6 User Filters exist. This feature can be used by clients configured to use the SSL protocol. This feature does not support IKEv2/IPsec protocol. We modified these screen: Monitoring > VPN > VPN Statistics > Sessions.
NAT support for IPv6	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6 (NAT64). Translating between IPv4 and IPv6 is not supported in transparent mode. We modified the following screens: Configuration > Firewall > Objects > Network Objects/Group Configuration > Firewall > NAT Rules
DHCPv6 relay	DHCP relay is supported for IPv6. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
OSPFv3	OSPFv3 routing is supported for IPv6. Note the following additional guidelines and limitations for OSPFv2 and OSPFv3: **Clustering** • OSPFv2 and OSPFv3 support clustering. • When clustering is configured, OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment. • When using individual interfaces, make sure that you establish the master and slave units as either OSPFv2 or OSPFv3 neighbors. • When using individual interfaces, OSPFv2 adjacencies can only be established between two contexts on a shared interface on the master unit. Configuring static neighbors is supported only on point-to-point links; therefore, only one neighbor statement is allowed on an interface. **Other** • OSPFv2 and OSPFv3 support multiple instances on an interface. • The ESP and AH protocol is supported for OSPFv3 authentication. • OSPFv3 supports Non-Payload Encryption. We introduced the following screens: Configuration > Device Setup > Routing > OSPFv3 > Setup Configuration > Device Setup > Routing > OSPFv3 > Interface Configuration > Device Setup > Routing > OSPFv3 > Redistribution Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix Configuration > Device Setup > Routing > OSPFv3 > Virtual Link Monitoring > Routing > OSPFv3 LSAs Monitoring > Routing > OSPFv3 Neighbors
Unified ACL for IPv4 and IPv6	ACLs now support IPv4 and IPv6 addresses. You can also specify a mix of IPv4 and IPv6 addresses for the source and destination. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. ACLs containing IPv6 addresses can be applied to clients configured to use the SSL protocol. This feature is not supported for the IKEv2/IPsec protocol. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Mixed IPv4 and IPv6 object groups	Previously, network object groups could only contain all IPv4 addresses or all IPv6 addresses. Now network object groups can support a mix of both IPv4 and IPv6 addresses. **Note** You cannot use a mixed object group for NAT. We modified the following screen: Configuration > Firewall > Objects > Network Objects/Groups.
Range of IPv6 addresses for a Network object	You can now configure a range of IPv6 addresses for a network object. We modified the following screen: Configuration > Firewall > Objects > Network Objects/Groups.
Inspection support for IPv6 and NAT64	We now support DNS inspection for IPv6 traffic. We also support translating between IPv4 and IPv6 for the following inspections: • DNS • FTP • HTTP • ICMP You can now also configure the service policy to generate a syslog message (767001) when unsupported inspections receive and drop IPv6 traffic. We modified the following screen: Configuration > Firewall > Service Policy Rules > Add Service Policy Rule Wizard - Service Policy.
Remote Access Features	
Clientless SSL VPN: Additional Support	We have added additional support for these browsers, operating systems, web technologies and applications: **Internet browser support**: Microsoft Internet Explorer 9, Firefox 4, 5, 6, 7, and 8 **Operating system support**: Mac OS X 10.7 **Web technology support**: HTML 5 **Application Support**: Sharepoint 2010
Clientless SSL VPN: Enhanced quality for rewriter engines	The clientless SSL VPN rewriter engines were significantly improved to provide better quality and efficacy. As a result, you can expect a better end-user experience for clientless SSL VPN users. We did not add or modify any ASDM screens for this feature. *Also available in 8.4(4.1).*

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Clientless SSL VPN: Citrix Mobile Receiver	This feature provides secure remote access for Citrix Receiver applications running on mobile devices to XenApp and XenDesktop VDI servers through the ASA. For the ASA to proxy Citrix Receiver to a Citrix Server, when users try to connect to Citrix virtualized resource, instead of providing the Citrix Server's address and credentials, users enter the ASA's SSL VPN IP address and credentials. We modified the following screen: Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policy > Edit > More Options > VDI Access > Add VDI Server.
Clientless SSL VPN: Enhanced Auto-sign-on	This feature improves support for web applications that require dynamic parameters for authentication. We modified the following screen: Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks.
Clientless SSL VPN: Clientless Java Rewriter Proxy Support	This feature provides proxy support for clientless Java plug-ins when a proxy is configured in client machines' browsers. We did not add or modify any ASDM screens for this feature.
Clientless SSL VPN: Remote File Explorer	The Remote File Explorer provides users with a way to browse the corporate network from their web browser. When users click the Remote File System icon on the Cisco SSL VPN portal page, an applet is launched on the user's system displaying the remote file system in a tree and folder view. We did not add or modify any ASDM screens for this feature.
Clientless SSL VPN: Server Certificate Validation	This feature enhances clientless SSL VPN support to enable SSL server certificate verification for remote HTTPS sites against a list of trusted CA certificates. We modified the following screen: Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool.
AnyConnect Performance Improvements	This feature improves throughput performance for AnyConnect TLS/DTLS traffic in multi-core platforms. It accelerates the SSL VPN datapath and provides customer-visible performance gains in AnyConnect, smart tunnels, and port forwarding. We modified the following screen: Configuration > Remote Access VPN > Advanced > Crypto Engine.
Custom Attributes	Custom attributes define and configure AnyConnect features that have not yet been added to ASDM. You add custom attributes to a group policy, and define values for those attributes. For AnyConnect 3.1, custom attributes are available to support AnyConnect Deferred Upgrade. Custom attributes can benefit AnyConnect clients configured for either IKEv2/IPsec or SSL protocols. A new screen was added: Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Custom Attributes.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Next Generation Encryption	The National Standards Association (NSA) specified a set of cryptographic algorithms that devices must support to meet U.S. federal standards for cryptographic strength. RFC 6379 defines the Suite B cryptographic suites. Because the collective set of algorithms defined as NSA Suite B are becoming a standard, the AnyConnect IPsec VPN (IKEv2 only) and public key infrastructure (PKI) subsystems now support them. The next generation encryption (NGE) includes a larger superset of this set adding cryptographic algorithms for IPsec V3 VPN, Diffie-Hellman Groups 14 and 24 for IKEv2, and RSA certificates with 4096 bit keys for DTLS and IKEv2. The following functionality is added to ASA to support the Suite B algorithms: • AES-GCM/GMAC support (128-, 192-, and 256-bit keys) – IKEv2 payload encryption and authentication – ESP packet encryption and authentication – Hardware supported only on multi-core platforms • SHA-2 support (256-, 384-, and 512-bit hashes) – ESP packet authentication – Hardware and software supported only on multi-core platforms • ECDH support (groups 19, 20, and 21) – IKEv2 key exchange – IKEv2 PFS – Software only supported on single- or multi-core platforms • ECDSA support (256-, 384-, and 521-bit elliptic curves) – IKEv2 user authentication – PKI certificate enrollment – PKI certificate generation and verification – Software only supported on single- or multi-core platforms New cryptographic algorithms are added for IPsecV3. **Note** Suite B algorithm support requires an AnyConnect Premium license for IKEv2 remote access connections, but Suite B usage for other connections or purposes (such as PKI) has no limitations. IPsecV3 has no licensing restrictions. We introduced or modified the following screens: Monitor > VPN > Sessions Monitor > VPN > Encryption Statistics Configuration > Site-to-Site VPN > Certificate Management > Identity Certificates Configuration > Site-to-Site VPN > Advanced > System Options Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > Crypto Maps
Support for VPN on the ASASM	The ASASM now supports all VPN features.

Table 1-4 New Features for ASA Version 9.0(1)/ASDM Version 7.0(1) (continued)

Feature	Description
Multiple Context Mode Features	
Site-to-Site VPN in multiple context mode	Site-to-site VPN tunnels are now supported in multiple context mode.
New resource type for site-to-site VPN tunnels	New resource types, vpn other and vpn burst other, were created to set the maximum number of site-to-site VPN tunnels in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class.
Dynamic routing in Security Contexts	EIGRP and OSPFv2 dynamic routing protocols are now supported in multiple context mode. OSPFv3, RIP, and multicast routing are not supported.
New resource type for routing table entries	A new resource class, routes, was created to set the maximum number of routing table entries in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class.
Mixed firewall mode support in multiple context mode	You can set the firewall mode independently for each security context in multiple context mode, so some can run in transparent mode while others run in routed mode. You cannot set the firewall mode in ASDM; you must use the command-line interface. *Also available in Version 8.5(1).*
Module Features	
ASA Services Module support on the Cisco 7600 switch	The Cisco 7600 series now supports the ASASM. For specific hardware and software requirements, see: http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html.
ASA 5585-X support for the ASA CX SSP-10 and -20	The ASA CX module lets you enforce security based on the complete context of a situation. This context includes the identity of the user (who), the application or website that the user is trying to access (what), the origin of the access attempt (where), the time of the attempted access (when), and the properties of the device used for the access (how). With the ASA CX module, you can extract the full context of a flow and enforce granular policies such as permitting access to Facebook but denying access to games on Facebook or permitting finance employees access to a sensitive enterprise database but denying the same to other employees. We introduced the following screens: Home > ASA CX Status Wizards > Startup Wizard > ASA CX Basic Configuration Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA CX Inspection *Also available in 8.4(4.1).*
ASA 5585-X Dual SSP support for the SSP-10 and SSP-20 (in addition to the SSP-40 and SSP-60); VPN support for Dual SSPs	The ASA 5585-X now supports dual SSPs using all SSP models (you can use two SSPs of the same level in the same chassis). VPN is now supported when using dual SSPs. We did not modify any screens.

How the ASA Services Module Works with the Switch

You can install the ASASM in the Catalyst 6500 series and Cisco 7600 series switches with Cisco IOS software on both the switch supervisor and the integrated MSFC.

Note The Catalyst Operating System (OS) is not supported.

The ASA runs its own operating system.

The switch includes a switching processor (the supervisor) and a router (the MSFC). Although you need the MSFC as part of your system, you do not have to use it. If you choose to do so, you can assign one or more VLAN interfaces to the MSFC. You can alternatively use external routers instead of the MSFC.

In single context mode, you can place the router in front of the firewall or behind the firewall (see Figure 1-1).

The location of the router depends entirely on the VLANs that you assign to it. For example, the router is behind the firewall in the example shown on the left side of Figure 1-1 because you assigned VLAN 201 to the inside interface of the ASASM. The router is in front of the firewall in the example shown on the right side of Figure 1-1 because you assigned VLAN 200 to the outside interface of the ASASM.

In the left-hand example, the MSFC or router routes between VLANs 201, 301, 302, and 303, and no inside traffic goes through the ASASM unless it is destined for the Internet. In the right-hand example, the ASASM processes and protects all traffic between the inside VLANs 201, 202, and 203.

Figure 1-1 MSFC/Router Placement

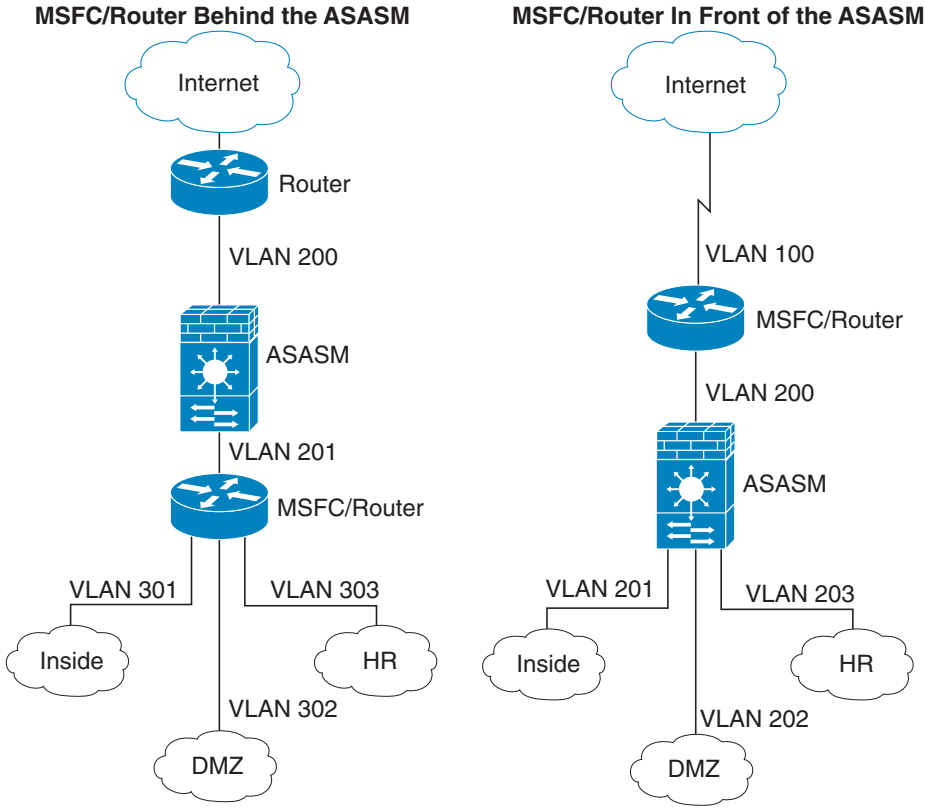

For multiple context mode, if you place the router behind the ASASM, you should only connect it to a single context. If you connect the router to multiple contexts, the router will route between the contexts, which might not be your intention. The typical scenario for multiple contexts is to use a router in front of all the contexts to route between the Internet and the switched networks (see Figure 1-2).

Figure 1-2 MSFC/Router Placement with Multiple Contexts

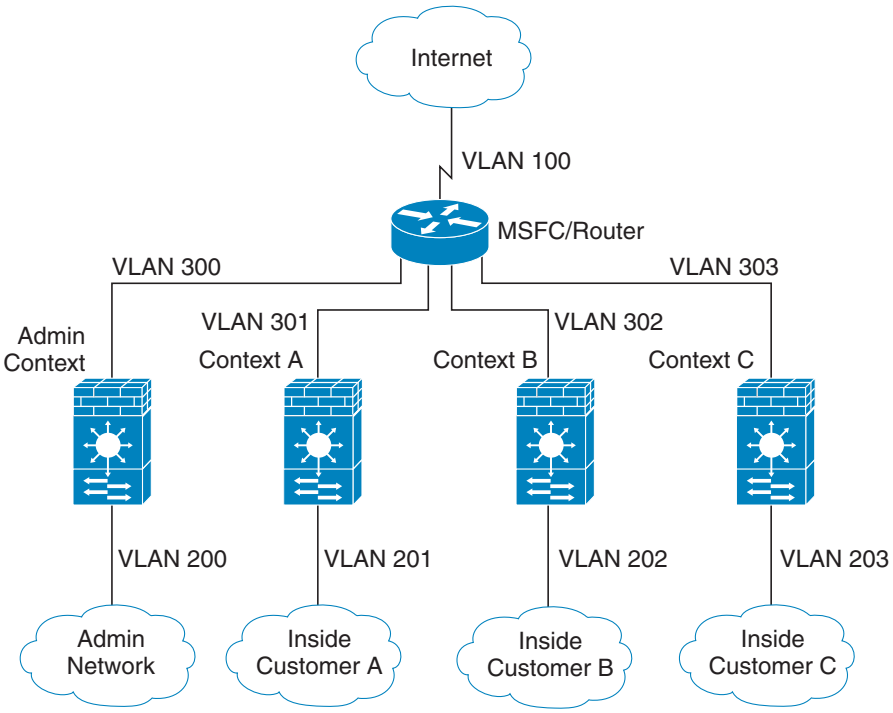

Firewall Functional Overview

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 ASA 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:

- Security Policy Overview, page 1-21
- Firewall Mode Overview, page 1-23
- Stateful Inspection Overview, page 1-24

Security Policy Overview

A security policy determines which traffic is allowed to pass through the firewall to access another network. By default, the ASA 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 Rules, page 1-21
- Applying NAT, page 1-21
- Protecting from IP Fragments, page 1-22
- Using AAA for Through Traffic, page 1-22
- Applying HTTP, HTTPS, or FTP Filtering, page 1-22
- Applying Application Inspection, page 1-22
- Sending Traffic to the IPS Module, page 1-22
- Sending Traffic to the Content Security and Control Module, page 1-22
- Applying QoS Policies, page 1-22
- Applying Connection Limits and TCP Normalization, page 1-23
- Enabling Threat Detection, page 1-23
- Enabling the Botnet Traffic Filter, page 1-23
- Configuring Cisco Unified Communications, page 1-23

Permitting or Denying Traffic with Access Rules

You can apply an access rule 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 ASA 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 ASA. 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 ASA 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 ASA 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 ASA to do a deep packet inspection.

Sending Traffic to the IPS Module

If your model supports the IPS module for intrusion prevention, then you can send traffic to the module for inspection. The IPS module 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 the documentation for your IPS module.

Sending Traffic to the Content Security and Control 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 ASA 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 ASA 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 ASA scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

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

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

Enabling 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 any suspicious activity. When you see syslog messages about the malware activity, you can take steps to isolate and disinfect the host.

Configuring Cisco Unified Communications

The Cisco ASA 5500 series is a strategic platform to provide proxy functions for unified communications deployments. The purpose of a proxy is to terminate and reoriginate connections between a client and server. The proxy delivers a range of security functions such as traffic inspection, protocol conformance, and policy control to ensure security for the internal network. An increasingly popular function of a proxy is to terminate encrypted connections in order to apply security policies while maintaining confidentiality of connections.

Firewall Mode Overview

The ASA runs in two different firewall modes:

- Routed

- Transparent

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

In transparent mode, the ASA acts like a "bump in the wire," or a "stealth firewall," and is not considered a router hop. The ASA 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 ASA 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.

Note The TCP state bypass feature allows you to customize the packet flow. See the "TCP State Bypass" section on page 71-3.

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

- Is this a new connection?

  If it is a new connection, the ASA 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"

  The ASA creates forward and reverse flows in the fast path for TCP traffic; the ASA also creates connection state information for connectionless protocols like UDP, ICMP (when you enable ICMP inspection), so that they can also use the fast path.

  Note For other IP protocols, like SCTP, the ASA does not create reverse path flows. As a result, ICMP error packets that refer to these connections are dropped.

  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 ASA 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

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 ASA uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The ASA 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 ASA invokes various standard protocols to accomplish these functions.

The ASA 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 ASA invokes various standard protocols to accomplish these functions.

Security Context Overview

You can partition a single ASA 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 ASA 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 ASA. 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.

ASA Clustering Overview

ASA Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

You perform all configuration (aside from the bootstrap configuration) on the master unit only; the configuration is then replicated to the member units.

Configuring the Switch for Use with the ASA Services Module

This chapter describes how to configure the Catalyst 6500 series or Cisco 7600 series switch for use with the ASASM. Before completing the procedures in this chapter, configure the basic properties of your switch, including assigning VLANs to switch ports, according to the documentation that came with your switch.

This chapter includes the following sections:

- Information About the Switch, page 2-1
- Guidelines and Limitations, page 2-2
- Verifying the Module Installation, page 2-3
- Assigning VLANs to the ASA Services Module, page 2-4
- Using the MSFC as a Directly Connected Router, page 2-5
- Configuring the Switch for ASA Failover, page 2-9
- Resetting the ASA Services Module, page 2-11
- Monitoring the ASA Services Module, page 2-11
- Feature History for the Switch for Use with the ASA Services Module, page 2-13

Information About the Switch

You can install the ASASM in the Catalyst 6500 series or Cisco 7600 series switches. The switch includes a switch (the supervisor engine) as well as a router (the MSFC). The connection between the ASASM and the switch is a single 20-GB interface.

The switch supports Cisco IOS software on both the switch supervisor engine and the integrated MSFC router.

Note The Catalyst operating system software is not supported.

The ASASM runs its own operating system.

Note Because the ASASM runs its own operating system, upgrading the Cisco IOS software does not affect the operation of the ASASM.

To view a matrix of hardware and software compatibility for the ASASM and Cisco IOS versions, see the *Cisco ASA 5500 Series Hardware and Software Compatibility*:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

Some ASASM features interact with Cisco IOS features. The following features involve Cisco IOS software:

- Virtual Switching System (VSS)—No ASASM configuration is required.
- Autostate—The supervisor informs the ASASM when the last interface on a given VLAN has gone down, which assists in determining whether or not a failover switch is required.
- Clearing entries in the supervisor MAC address table on a failover switch—No ASASM configuration is required.
- Version compatibility—The ASASM will be automatically powered down if the supervisor/ASASM version compatibility matrix check fails.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

VLAN Guidelines and Limitations

- Use VLAN IDs 2 to 1000 and from 1025 to 4094.
- Routed ports and WAN ports consume internal VLANs, so it is possible that VLANs in the 1020-1100 range might already be in use.
- You cannot use reserved VLANs.
- You cannot use VLAN 1.
- If you are using ASASM failover within the same switch chassis, do not assign the VLAN(s) that you are reserving for failover and stateful communications to a switch port. However, if you are using failover between chassis, you must include the VLANs in the trunk port between the chassis.
- If you do not add the VLANs to the switch before you assign them to the ASASM, the VLANs are stored in the supervisor engine database and are sent to the ASASM as soon as they are added to the switch.
- You can configure a VLAN in the ASASM configuration before it has been assigned on the switch. Note that when the switch sends the VLAN to the ASASM, the VLAN defaults to be administratively up on the ASASM, regardless of whether the you shut them down in the ASASM configuration. You need to shut them down again in this case.

SPAN Reflector Guidelines

In Cisco IOS software Version 12.2SXJ1 and earlier, for each ASASM in a switch, the SPAN reflector feature is enabled. This feature allows multicast traffic (and other traffic that requires a central rewrite engine) to be switched when coming from the ASASM. The SPAN reflector feature uses one SPAN session. To disable this feature, enter the following command:

```
Router(config)# no monitor session servicemodule
```

Verifying the Module Installation

To verify that the switch acknowledges the ASASM and has brought it online, enter the following command.

Detailed Steps

Command	Purpose
show module [**switch** {**1** |**2**}] [*mod-num* | **all**]	Displays module information. For a switch in a VSS, enter the **switch** keyword.
Example: `Router# show module 1`	Ensure that the Status column shows "Ok" for the ASASM.

Examples

The following is sample output from the **show module** command:

```
Router# show module
Mod Ports Card Type                                Model              Serial No.
--- ----- -------------------------------------- ------------------ -----------
 2    3  ASA Service Module                     WS-SVC-ASA-SM1     SAD143502E8
 4    3  ASA Service Module                     WS-SVC-ASA-SM1     SAD135101Z9
 5    5  Supervisor Engine 720 10GE (Active)    VS-S720-10G        SAL12426KB1
 6   16  CEF720 16 port 10GE                    WS-X6716-10GE      SAL1442WZD1

Mod MAC addresses                       Hw    Fw           Sw           Status
--- ----------------------------------- ------ ------------ ------------ -------
 2  0022.bdd4.016f to 0022.bdd4.017e   0.201 12.2(2010080 12.2(2010121 Ok
 4  0022.bdd3.f64e to 0022.bdd3.f655   0.109 12.2(2010080 12.2(2010121 PwrDown
 5  0019.e8bb.7b0c to 0019.e8bb.7b13   2.0   8.5(2)       12.2(2010121 Ok
 6  f866.f220.5760 to f866.f220.576f   1.0   12.2(18r)S1  12.2(2010121 Ok

Mod  Sub-Module                  Model              Serial       Hw     Status
---- --------------------------- ------------------ ------------ ------- -------
2/0 ASA Application Processor   SVC-APP-PROC-1     SAD1436015D  0.202  Other
4/0 ASA Application Processor   SVC-APP-INT-1      SAD141002AK  0.106  PwrDown
 5  Policy Feature Card 3       VS-F6K-PFC3C       SAL12437BM2  1.0    Ok
 5  MSFC3 Daughterboard         VS-F6K-MSFC3       SAL12426DE3  1.0    Ok
 6  Distributed Forwarding Card WS-F6700-DFC3C     SAL1443XRDC  1.4    Ok

Base PID:
Mod  Model            Serial No.
---- ------------      ----------
 2 WS-SVC-APP-HW-1    SAD143502E8
 4 TRIFECTA           SAD135101Z9

Mod  Online Diag Status
---- -------------------
 2  Pass
2/0 Not Applicable
 4  Not Applicable
4/0 Not Applicable
 5  Pass
 6  Pass
```

Assigning VLANs to the ASA Services Module

This section describes how to assign VLANs to the ASASM. The ASASM does not include any external physical interfaces. Instead, it uses VLAN interfaces. Assigning VLANs to the ASASM is similar to assigning a VLAN to a switch port; the ASASM includes an internal interface to the Switch Fabric Module (if present) or the shared bus.

Prerequisites

See the switch documentation for information about adding VLANs to the switch and assigning them to switch ports.

Guidelines

- You can assign up to 16 firewall VLAN groups to each ASASM. (You can create more than 16 VLAN groups in Cisco IOS software, but only 16 can be assigned per ASASM.) For example, you can assign all the VLANs to one group; or you can create an inside group and an outside group; or you can create a group for each customer.
- There is no limit on the number of VLANs per group, but the ASASM can only use VLANs up to the ASASM system limit (see the ASASM licensing documentation for more information).
- You cannot assign the same VLAN to multiple firewall groups.
- You can assign a single firewall group to multiple ASASMs. VLANs that you want to assign to multiple ASASMs, for example, can reside in a separate group from VLANs that are unique to each ASASM.
- See the "VLAN Guidelines and Limitations" section on page 2-2.

Detailed Steps

	Command	Purpose
Step 1	**firewall vlan-group** *firewall_group vlan_range* **Example:** `Router(config)# firewall vlan-group 50 55-57`	Assigns VLANs to a firewall group. The *firewall_group* argument is an integer. The *vlan_range* argument can be one or more VLANs (2 to 1000 and from 1025 to 4094) identified in one of the following ways: • A single number (*n*) • A range (*n-x*) Separate numbers or ranges by commas, as shown in the following example: **5,7-10,13,45-100**
Step 2	**firewall** [**switch** {**1** |**2**}] **module** *slot* **vlan-group** *firewall_group* **Example:** `Router(config)# firewall module 5 vlan-group 50,52`	Assigns the firewall groups to the ASASM. For a switch in a VSS, enter the **switch** argument. To view the slots where the ASASM is installed, enter the **show module** command. The *firewall_group* argument is one or more group numbers, which can be one of the following: • A single number (*n*) • A range (*n-x*) Separate numbers or ranges by commas, as shown in the following example: **5,7-10**

Examples

The following example shows how to create three firewall VLAN groups: one for each ASASM, and one that includes VLANs assigned to both ASASMs:

```
Router(config)# firewall vlan-group 50 55-57
Router(config)# firewall vlan-group 51 70-85
Router(config)# firewall vlan-group 52 100
Router(config)# firewall module 5 vlan-group 50,52
Router(config)# firewall module 8 vlan-group 51,52
```

Using the MSFC as a Directly Connected Router

If you want to use the MSFC as a directly connected router (for example, as the default gateway connected to the ASASM outside interface), then add an ASASM VLAN interface to the MSFC as a switched virtual interface (SVI).

This section includes the following topics:

• Information About SVIs, page 2-6
• Configuring SVIs, page 2-8

Information About SVIs

For security reasons, by default, you can configure one SVI between the MSFC and the ASASM; you can enable multiple SVIs, but be sure you do not misconfigure your network.

For example, with multiple SVIs, you could accidentally allow traffic to pass around the ASASM by assigning both the inside and outside VLANs to the MSFC. (See Figure 2-1.)

Figure 2-1 Multiple SVI Misconfiguration

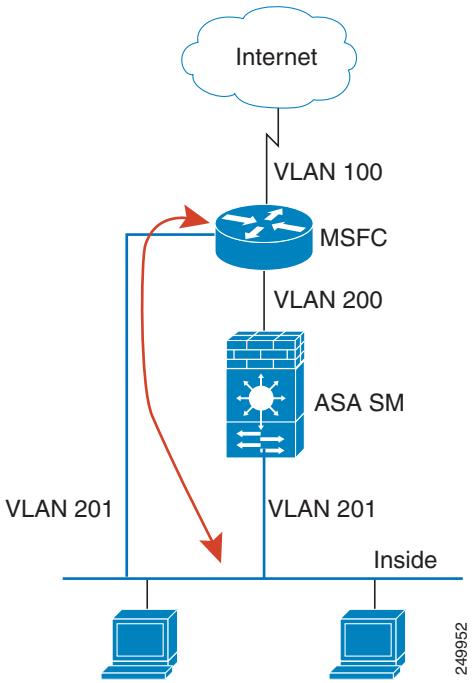

You might need to bypass the ASASM in some network scenarios. Figure 2-2 shows an IPX host on the same Ethernet segment as IP hosts. Because the ASASM in routed firewall mode only handles IP traffic and drops other protocol traffic like IPX (transparent firewall mode can optionally allow non-IP traffic), you might want to bypass the ASASM for IPX traffic. Make sure that you configure the MSFC with an access list that allows only IPX traffic to pass on VLAN 201.

Figure 2-2 Multiple SVIs for IPX

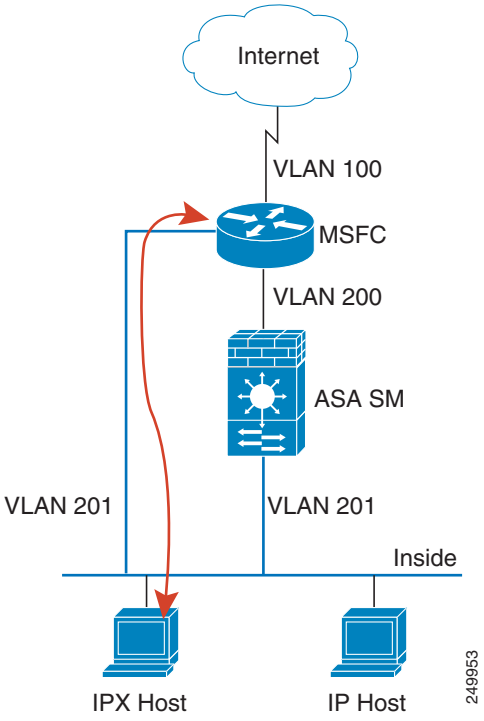

For transparent firewalls in multiple context mode, you need to use multiple SVIs because each context requires a unique VLAN on its outside interface (see Figure 2-3). You might also choose to use multiple SVIs in routed mode so that you do not have to share a single VLAN for the outside interface.

Figure 2-3 Multiple SVIs in Multiple Context Mode

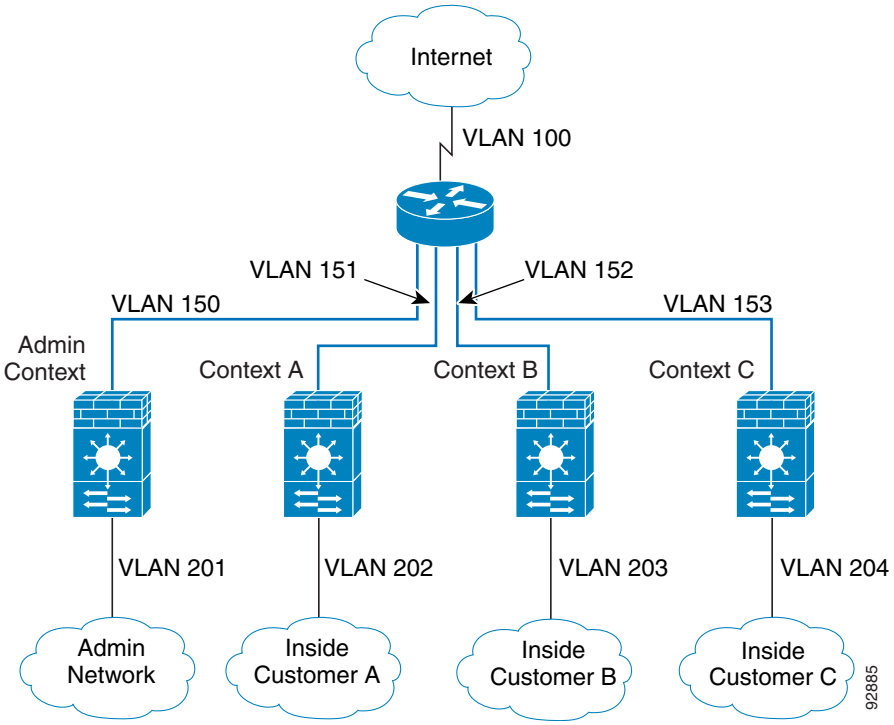

Configuring SVIs

To add an SVI to the MSFC, perform the following steps.

Detailed Steps

	Command	Purpose
Step 1	(Optional) **`firewall multiple-vlan-interfaces`** **Example:** `Router(config)# firewall multiple-vlan-interfaces`	Allows you to add more than one SVI to the ASASM.
Step 2	**`interface vlan`** *`vlan_number`* **Example:** `Router(config)# interface vlan 55`	Adds a VLAN interface to the MSFC.

	Command	Purpose
Step 3	**ip address** *address mask* **Example:** `Router(config-if)# ip address 10.1.1.1 255.255.255.0`	Sets the IP address for this interface on the MSFC.
Step 4	**no shutdown** **Example:** `Router(config-if)# no shutdown`	Enables the interface.

Examples

The following example shows a typical configuration with multiple SVIs:

```
Router(config)# firewall vlan-group 50 55-57
Router(config)# firewall vlan-group 51 70-85
Router(config)# firewall module 8 vlan-group 50-51
Router(config)# firewall multiple-vlan-interfaces
Router(config)# interface vlan 55
Router(config-if)# ip address 10.1.1.1 255.255.255.0
Router(config-if)# no shutdown
Router(config-if)# interface vlan 56
Router(config-if)# ip address 10.1.2.1 255.255.255.0
Router(config-if)# no shutdown
Router(config-if)# end
Router#
```

Configuring the Switch for ASA Failover

This section includes the following topics:

- Assigning VLANs to the Secondary ASA Services Module, page 2-10
- Adding a Trunk Between a Primary Switch and Secondary Switch, page 2-10
- Ensuring Compatibility with Transparent Firewall Mode, page 2-10
- Enabling Autostate Messaging for Rapid Link Failure Detection, page 2-10

Assigning VLANs to the Secondary ASA Services Module

Because both units require the same access to the inside and outside networks, you must assign the same VLANs to both ASASMs on the switch(es). See the "Assigning VLANs to the Secondary ASA Services Module" section on page 2-10.

Adding a Trunk Between a Primary Switch and Secondary Switch

If you are using inter-switch failover, then you should configure an 802.1Q VLAN trunk between the two switches to carry the failover and state links. The trunk should have QoS enabled so that failover VLAN packets, which have a CoS value of 5 (higher priority), are treated with higher priority in these ports.

To configure the EtherChannel and trunk, see the documentation for your switch.

Ensuring Compatibility with Transparent Firewall Mode

To avoid loops when you use failover in transparent mode, use switch software that supports BPDU forwarding. Do not enable LoopGuard globally on the switch if the ASASM is in transparent mode. LoopGuard is automatically applied to the internal EtherChannel between the switch and the ASASM, so after a failover and a failback, LoopGuard causes the secondary unit to be disconnected because the EtherChannel goes into the err-disable state.

Enabling Autostate Messaging for Rapid Link Failure Detection

The supervisor engine can send autostate messages to the ASASM about the status of physical interfaces associated with ASASM VLANs. For example, when all physical interfaces associated with a VLAN go down, the autostate message tells the ASASM that the VLAN is down. This information lets the ASASM declare the VLAN as down, bypassing the interface monitoring tests normally required for determining which side suffered a link failure. Autostate messaging provides a dramatic improvement in the time the ASASM takes to detect a link failure (a few milliseconds as compared to up to 45 seconds without autostate support).

The switch supervisor sends an autostate message to the ASASM when:

- The last interface belonging to a VLAN goes down.
- The first interface belonging to a VLAN comes up.

Detailed Steps

Command	Purpose
firewall autostate **Example:** `Router(config)# firewall autostate`	Enables autostate messaging in Cisco IOS software. Autostate messaging is disabled by default.

Resetting the ASA Services Module

This section describes how to reset the ASASM. You might need to reset the ASASM if you cannot reach it through the CLI or an external Telnet session. The reset process might take several minutes.

Detailed Steps

Command	Purpose
hw-module [**switch** {**1** |**2**}] **module** *slot* **reset** **Example:** `Router# hw-module module 9 reset`	Resets the ASASM. For a switch in a VSS, enter the **switch** argument. The *slot* argument indicates the slot number in which the module is installed. To view the slots where the ASASM is installed, enter the **show module** command. **Note** To reset the ASASM when you are already logged in to it. enter either the **reload** or **reboot** command.

Examples

The following is sample output from the **hw-module module reset** command:

```
Router# hw-module module 9 reset

Proceed with reload of module? [confirm] y
% reset issued for module 9

Router#
00:26:55:%SNMP-5-MODULETRAP:Module 9 [Down] Trap
00:26:55:SP:The PC in slot 8 is shutting down. Please wait ...
```

Monitoring the ASA Services Module

To monitor the ASA, enter one of the following commands:

Command	Purpose
show firewall module [*mod-num*] **state**	Verifies the state of the ASA.
show firewall module [*mod-num*] **traffic**	Verifies that traffic is flowing through the ASA.
show firewall module [*mod-num*] **version**	Shows the software version of the ASA.
show firewall multiple-vlan-interfaces	Indicates the status of multiple VLAN interfaces (enabled or disabled).
show firewall vlan-group	Displays all configured VLAN groups.
show interface vlan	Displays the status and information about the configured VLAN interface.

Examples

The following is sample output from the **show firewall module** [*mod-num*] **state** command:

```
Router> show firewall module 11 state
Firewall module 11:
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 3,6,7,20-24,40,59,85,87-89,99-115,150,188-191,200,250,
     501-505,913,972
Pruning VLANs Enabled: 2-1001
Vlans allowed on trunk:
Vlans allowed and active in management domain:
Vlans in spanning tree forwarding state and not pruned:
```

The following is sample output from the **show firewall module** [*mod-num*] **traffic** command:

```
Router> show firewall module 11 traffic
Firewall module 11:

Specified interface is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 0014.1cd5.bef6 (bia 0014.1cd5.bef6)
  MTU 1500 bytes, BW 6000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Full-duplex, 1000Mb/s, media type is unknown
  input flow-control is on, output flow-control is on
  Members in this channel: Gi11/1 Gi11/2 Gi11/3 Gi11/4 Gi11/5 Gi11/6
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queuing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 10000 bits/sec, 17 packets/sec
     8709 packets input, 845553 bytes, 0 no buffer
     Received 745 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     18652077 packets output, 1480488712 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

The following is sample output from the **show firewall multiple-vlan-interfaces** command:

```
Router# show firewall multiple-vlan-interfaces
Multiple firewall vlan interfaces feature is enabled
```

The following is sample output from the **show firewall module** command:

```
Router# show firewall module
Module Vlan-groups
  5    50,52
  8    51,52
```

The following is sample output from the **show firewall module** [*mod-num*] **version** command:

```
Router# show firewall module 2 version
ASA Service Module 2:

Sw Version: 100.7(8)19
```

The following is sample output from the **show firewall vlan-group** command:

```
Router# show firewall vlan-group
Group vlans
----- ------
   50 55-57
   51 70-85
   52 100
```

The following is sample output from the **show interface vlan** command:

```
Router# show interface vlan 55
Vlan55 is up, line protocol is up
  Hardware is EtherSVI, address is 0008.20de.45ca (bia 0008.20de.45ca)
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  ARP type:ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:08, output hang never
  Last clearing of "show interface" counters never
  Input queue:0/75/0/0 (size/max/drops/flushes); Total output drops:0
  Queueing strategy:fifo
  Output queue :0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  L2 Switched:ucast:196 pkt, 13328 bytes - mcast:4 pkt, 256 bytes
  L3 in Switched:ucast:0 pkt, 0 bytes - mcast:0 pkt, 0 bytes mcast
  L3 out Switched:ucast:0 pkt, 0 bytes
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     4 packets output, 256 bytes, 0 underruns
     0 output errors, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

Feature History for the Switch for Use with the ASA Services Module

Table 2-1 lists each feature change and the platform release in which it was implemented.

Table 2-1 Feature History for the Switch for Use with the ASASM

Feature Name	Platform Releases	Feature Information
ASA Services Module support on the Cisco Catalyst 6500 switch	8.5(1)	The ASASM is a high-performance security services module for the Catalyst 6500 series switch, which you configure according to the procedures in this chapter. We introduced or modified the following commands: **firewall transparent, mac address auto, firewall autostate** (IOS), **interface vlan**.
ASA Services Module support on the Cisco 7600 switch	9.0(1)	The Cisco 7600 series now supports the ASASM.

Getting Started

This chapter describes how to get started with your ASA. This chapter includes the following sections:

- Accessing the Appliance Command-Line Interface, page 3-1
- Accessing the ASA Services Module Command-Line Interface, page 3-2
- Configuring ASDM Access for Appliances, page 3-6
- Configuring ASDM Access for the ASA Services Module, page 3-11
- Starting ASDM, page 3-13
- Factory Default Configurations, page 3-17
- Getting Started with the Configuration, page 3-24
- Using the Command Line Interface Tool in ASDM, page 3-25
- Applying Configuration Changes to Connections, page 3-27

Accessing the Appliance Command-Line Interface

In some cases, you may need to use the CLI to configure basic settings for ASDM access. See the "Configuring ASDM Access for Appliances" section on page 3-6 to determine if you need to use the CLI.

For initial configuration, access the CLI directly from the console port. Later, you can configure remote access using Telnet or SSH according to Chapter 52, "Configuring Management Access." If your system is already in multiple context mode, then accessing the console port places you in the system execution space. See Chapter 9, "Configuring Multiple Context Mode," for more information about multiple context mode.

Detailed Steps

Step 1 Connect a PC to the console port using the provided console cable, and connect to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, no flow control.

See the hardware guide for your ASA for more information about the console cable.

Step 2 Press the **Enter** key to see the following prompt:

hostname>

This prompt indicates that you are in user EXEC mode. Only basic commands are available from user EXEC mode.

Step 3 To access privileged EXEC mode, enter the following command:

```
hostname> enable
```

The following prompt appears:

```
Password:
```

All non-configuration commands are available in privileged EXEC mode. You can also enter configuration mode from privileged EXEC mode.

Step 4 Enter the enable password at the prompt.

By default, the password is blank, and you can press the **Enter** key to continue. See the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1 to change the enable password.

The prompt changes to:

```
hostname#
```

To exit privileged mode, enter the **disable**, **exit**, or **quit** command.

Step 5 To access global configuration mode, enter the following command:

```
hostname# configure terminal
```

The prompt changes to the following:

```
hostname(config)#
```

You can begin to configure the ASA from global configuration mode. To exit global configuration mode, enter the **exit**, **quit**, or **end** command.

Accessing the ASA Services Module Command-Line Interface

For initial configuration, access the command-line interface by connecting to the switch (either to the console port or remotely using Telnet or SSH) and then connecting to the ASASM. The ASASM does not include a factory default configuration, so you must perform some configuration at the CLI before you can access it using ASDM. This section describes how to access the ASASM CLI, and includes the following sections:

- Logging Into the ASA Services Module, page 3-2
- Logging Out of a Console Session, page 3-4
- Logging Out of a Telnet Session, page 3-6

Logging Into the ASA Services Module

For initial configuration, access the command-line interface by connecting to the switch (either to the switch console port or remotely using Telnet or SSH) and then connecting to the ASASM.

If your system is already in multiple context mode, then accessing the ASASM from the switch places you in the system execution space. See Chapter 9, "Configuring Multiple Context Mode," for more information about multiple context mode.

Later, you can configure remote access directly to the ASASM using Telnet or SSH according to the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 52-1.

This section includes the following topics:

- Information About Connection Methods, page 3-3
- Logging In, page 3-3

Information About Connection Methods

From the switch CLI, you can use two methods to connect to the ASASM:

- Telnet connection—Using the **session** command, you create a Telnet connection to the ASASM.

  Benefits include:

  - You can have multiple sessions to the ASASM at the same time.
  - The Telnet session is a fast connection.

  Limitations include:

  - The Telnet session is terminated when the ASASM reloads, and can time out.
  - You cannot access the ASASM until it completely loads; you cannot access ROMMON.

- Virtual console connection—Using the **service-module session** command, you create a virtual console connection to the ASASM, with all the benefits and limitations of an actual console connection.

  Benefits include:

  - The connection is persistent across reloads and does not time out.
  - You can stay connected through ASASM reloads and view startup messages.
  - You can access ROMMON if the ASASM cannot load the image.

  Limitations include:

  - The connection is slow (9600 baud).
  - You can only have one console connection active at a time.
  - You cannot use this command in conjunction with a terminal server where **Ctrl-Shift-6**, **x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6**, **x** is also the sequence to escape the ASASM console and return to the switch prompt. Therefore, if you try to exit the ASASM console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the switch, the ASASM console session is still active; you can never exit to the switch prompt. You must use a direct serial connection to return the console to the switch prompt. In this case, either change the terminal server or switch escape character in Cisco IOS, or use the Telnet **session** command instead.

Note Because of the persistence of the console connection, if you do not properly log out of the ASASM, the connection may exist longer than intended. If someone else wants to log in, they will need to kill the existing connection. See the "Logging Out of a Console Session" section on page 3-4 for more information.

Logging In

Perform the following steps to log into the ASASM and access global configuration mode.

Detailed Steps

	Command	Purpose
Step 1	From the switch, perform one of the following:	

	session [**switch** {**1** |**2**}] **slot** *number* **processor 1**
	
	You are prompted for the login password:
	`hostname passwd:`
	
	Example:
	`Router# session slot number processor 1`
	`hostname passwd: cisco`
	`hostname>`
		From the switch CLI, enter this command to Telnet to the ASASM over the backplane.
		
		For a switch in a VSS, enter the **switch** argument.
		
		Note The **session** *slot* **processor 0** command, which is supported on other services modules, is not supported on the ASASM; the ASASM does not have a processor 0.
		
		To view the module slot numbers, enter the **show module** command at the switch prompt.
		
		Enter the login password to the ASASM. By default, the password is **cisco**.
		
		You access user EXEC mode.

	service-module session [**switch** {**1** | **2**}] **slot** *number*
	
	Example:
	`Router# service-module session slot 3`
	`hostname>`
		From the switch CLI, enter this command to gain console access to the ASASM.
		
		For a switch in a VSS, enter the **switch** argument.
		
		To view the module slot numbers, enter the **show module** command at the switch prompt.
		
		You access user EXEC mode.

Step 2	**enable**
	
	Example:
	`hostname> enable`
	`Password:`
	`hostname#`
		Accesses privileged EXEC mode, which is the highest privilege level.
		
		Enter the enable password at the prompt. By default, the password is blank. To change the enable password, see the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1.
		
		To exit privileged EXEC mode, enter the **disable**, **exit**, or **quit** command.

Step 3	**configure terminal**
	
	Example:
	`hostname# configure terminal`
	`hostname(config)#`
		Accesses global configuration mode.
		
		To exit global configuration mode, enter the **disable**, **exit**, or **quit** command.

Logging Out of a Console Session

This section includes the following topics:

- Logging Out, page 3-5
- Killing an Active Console Connection, page 3-5

Logging Out

If you do not log out of the ASASM, the console connection persists; there is no timeout. To end the ASASM console session and access the switch CLI, perform the following steps.

To kill another user's active connection, which may have been unintentionally left open, see the "Killing an Active Console Connection" section on page 3-5.

Detailed Steps

Step 1 To return to the switch CLI, type the following:

Ctrl-Shift-6, **x**

You return to the switch prompt:

```
asasm# [Ctrl-Shift-6, x]
Router#
```

Note Shift-6 on US and UK keyboards issues the caret (^) character. If you have a different keyboard and cannot issue the caret (^) character as a standalone character, you can temporarily or permanently change the escape character to a different character. In Cisco IOS, before you session to the ASASM, use the **terminal escape-character** *ascii_number* command (to change temporarily) or the **default escape-character** *ascii_number* command (to change permanently). For example, to temporarily change the sequence to **Ctrl-w**, **x**, enter **terminal escape-character 23**. The next time you log into the switch, the escape character reverts back to the default.

Killing an Active Console Connection

Because of the persistence of a console connection, if you do not properly log out of the ASASM, the connection may exist longer than intended. If someone else wants to log in, they will need to kill the existing connection.

Detailed Steps

Step 1 From the switch CLI, show the connected users using the **show users** command. A console user is called "con". The Host address shown is 127.0.0.*slot*0, where *slot* is the slot number of the module.

```
Router# show users
```

For example, the following command output shows a user "con" on line 0 on a module in slot 2:

```
Router# show users
Line       User       Host(s)              Idle       Location
*  0       con 0     127.0.0.20            00:00:02
```

Step 2 To clear the line with the console connection, enter the following command:

```
Router# clear line number
```

For example:

```
Router# clear line 0
```

Logging Out of a Telnet Session

To end the Telnet session and access the switch CLI, perform the following steps.

Detailed Steps

Step 1 To return to the switch CLI, type **exit** from the ASASM privileged or user EXEC mode. If you are in a configuration mode, enter **exit** repeatedly until you exit the Telnet session.

You return to the switch prompt:

```
asasm# exit
Router#
```

Note You can alternatively escape the Telnet session using the escape sequence **Ctrl-Shift-6**, **x**; this escape sequence lets you resume the Telnet session by pressing the **Enter** key at the switch prompt. To disconnect your Telnet session from the switch, enter **disconnect** at the switch CLI. If you do not disconnect the session, it will eventually time out according to the ASASM configuration.

Configuring ASDM Access for Appliances

ASDM access requires some minimal configuration so you can communicate over the network with a management interface. This section includes the following topics:

- Accessing ASDM Using the Factory Default Configuration, page 3-6
- Accessing ASDM Using a Non-Default Configuration (ASA 5505), page 3-7
- Accessing ASDM Using a Non-Default Configuration (ASA 5510 and Higher), page 3-9

Accessing ASDM Using the Factory Default Configuration

With a factory default configuration (see the "Factory Default Configurations" section on page 3-17), ASDM connectivity is pre-configured with default network settings. Connect to ASDM using the following interface and network settings:

- The management interface depends on your model:
  - 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.
- The default management address is 192.168.1.1.

- The clients allowed to access ASDM must be on the 192.168.1.0/24 network. The default configuration enables DHCP so your management station can be assigned an IP address in this range. To allow other client IP addresses to access ASDM, see the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 52-1.

To launch ASDM, see the "Starting ASDM" section on page 3-13.

Note To change to multiple context mode, see the "Enabling or Disabling Multiple Context Mode" section on page 9-16. After changing to multiple context mode, you can access ASDM from the admin context using the network settings above.

Accessing ASDM Using a Non-Default Configuration (ASA 5505)

If you do not have a factory default configuration, want to change the configuration, or want to change to transparent firewall mode, perform the following steps. See also the sample configurations in the "ASA 5505 Default Configuration" section on page 3-20.

Prerequisites

Access the CLI according to the "Accessing the Appliance Command-Line Interface" section on page 3-1.

Detailed Steps

	Command	Purpose
Step 1	(Optional) **firewall transparent** **Example:** `hostname(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Configuring the Firewall Mode" section on page 6-1 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode:	

	Routed mode:
	interface vlan *number*
	ip address *ip_address* [*mask*]
	nameif *name*
	security-level *level*
	
	Example:
	`hostname(config)# interface vlan 1`
	`hostname(config-if)# ip address 192.168.1.1 255.255.255.0`
	`hostname(config-if)# nameif inside`
	`hostname(config-if)# security-level 100`
		Configures an interface in routed mode. The **security-level** is a number between 1 and 100, where 100 is the most secure.

	Command	Purpose

	Transparent mode:
	interface bvi *number*
	  ip address *ip_address* [*mask*]
	
	interface vlan *number*
	  bridge-group *bvi_number*
	  nameif *name*
	  security-level *level*
	
	Example:
	`hostname(config)# interface bvi 1`
	`hostname(config-if)# ip address 192.168.1.1 255.255.255.0`
	
	`hostname(config)# interface vlan 1`
	`hostname(config-if)# bridge-group 1`
	`hostname(config-if)# nameif inside`
	`hostname(config-if)# security-level 100`
		Configures a bridge virtual interface and assigns a management VLAN to the bridge group. The **security-level** is a number between 1 and 100, where 100 is the most secure.

Step 3	**interface ethernet 0/***n*
	  switchport access vlan *number*
	  no shutdown
	
	Example:
	`hostname(config)# interface ethernet 0/1`
	`hostname(config-if)# switchport access vlan 1`
	`hostname(config-if)# no shutdown`
		Enables the management switchport and assigns it to the management VLAN.

Step 4	**dhcpd address** *ip_address-ip_address interface_name*
	dhcpd enable *interface_name*
	
	Example:
	`hostname(config)# dhcpd address 192.168.1.5-192.168.1.254 inside`
	`hostname(config)# dhcpd enable inside`
		Enables DHCP for the management host on the management interface network. Make sure you do not include the management address in the range.
		
		Note By default, the IPS module, if installed, uses 192.168.1.2 for its internal management address, so be sure not to use this address in the DHCP range. You can later change the IPS module management address using the ASA if required.

Step 5	**http server enable** **Example:** `hostname(config)# http server enable`	Enables the HTTP server for ASDM.
Step 6	**http** *ip_address mask interface_name* **Example:** `hostname(config)# http 192.168.1.0 255.255.255.0 inside`	Allows the management host to access ASDM.
Step 7	**write memory** **Example:** `hostname(config)# write memory`	Saves the configuration.
Step 8	To launch ASDM, see the "Starting ASDM" section on page 3-13.	

Examples

The following configuration converts the firewall mode to transparent mode, configures the VLAN 1 interface and assigns it to BVI 1, enables a switchport, and enables ASDM for a management host:

```
firewall transparent
interface bvi 1
    ip address 192.168.1.1 255.255.255.0
interface vlan 1
    bridge-group 1
    nameif inside
    security-level 100
interface ethernet 0/1
    switchport access vlan 1
    no shutdown
dhcpd address 192.168.1.5-192.168.1.254 inside
dhcpd enable inside
http server enable
http 192.168.1.0 255.255.255.0 inside
```

Accessing ASDM Using a Non-Default Configuration (ASA 5510 and Higher)

If you do not have a factory default configuration, or want to change the firewall or context mode, perform the following steps.

Prerequisites

Access the CLI according to the "Accessing the Appliance Command-Line Interface" section on page 3-1.

Detailed Steps

	Command	Purpose
Step 1	(Optional) **firewall transparent** **Example:** `hostname(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Configuring the Firewall Mode" section on page 6-1 for more information.

Step 2	**interface management 0/0**
	ip address *ip_address mask*
	nameif *name*
	security-level *number*
	no shutdown
	
	Example:
	`hostname(config)# interface management 0/0`
	`hostname(config-if)# ip address 192.168.1.1 255.255.255.0`
	`hostname(config-if)# nameif management`
	`hostname(config-if)# security-level 100`
	`hostname(config-if)# no shutdown`
		Configures the Management 0/0 interface. The **security-level** is a number between 1 and 100, where 100 is the most secure.

	Command	Purpose

Step 3	(For directly-connected management hosts)
	dhcpd address *ip_address-ip_address interface_name*
	dhcpd enable *interface_name*
	
	Example:
	`hostname(config)# dhcpd address 192.168.1.2-192.168.1.254 management`
	`hostname(config)# dhcpd enable management`
		Enables DHCP for the management host on the management interface network. Make sure you do not include the Management 0/0 address in the range.

Step 4	(For remote management hosts) **route** *management_ifc management_host_ip mask gateway_ip* **1** **Example:** `hostname(config)# route management 10.1.1.0 255.255.255.0 192.168.1.50`	Configures a route to the management hosts.
Step 5	**http server enable** **Example:** `hostname(config)# http server enable`	Enables the HTTP server for ASDM.
Step 6	**http** *ip_address mask interface_name* **Example:** `hostname(config)# http 192.168.1.0 255.255.255.0 management`	Allows the management host to access ASDM.
Step 7	**write memory** **Example:** `hostname(config)# write memory`	Saves the configuration.
Step 8	(Optional) **mode multiple** **Example:** `hostname(config)# mode multiple`	Sets the mode to multiple mode. When prompted, confirm that you want to convert the existing configuration to be the admin context. You are then prompted to reload the ASASM. See Chapter 9, "Configuring Multiple Context Mode," for more information.
Step 9	To launch ASDM, see the "Starting ASDM" section on page 3-13.	

Examples

The following configuration converts the firewall mode to transparent mode, configures the Management 0/0 interface, and enables ASDM for a management host:

```
firewall transparent
interface management 0/0
    ip address 192.168.1.1 255.255.255.0
    nameif management
    security-level 100
```

```
    no shutdown
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd enable management
http server enable
http 192.168.1.0 255.255.255.0 management
```

Configuring ASDM Access for the ASA Services Module

Because the ASASM does not have physical interfaces, it does not come pre-configured for ASDM access; you must configure ASDM access using the CLI on the ASASM. To configure the ASASM for ASDM access, perform the following steps.

Prerequisites

- Assign a VLAN interface to the ASASM according to the "Assigning VLANs to the ASA Services Module" section on page 2-4.
- Connect to the ASASM and access global configuration mode according to the "Accessing the ASA Services Module Command-Line Interface" section on page 3-2.

Detailed Steps

	Command	Purpose
Step 1	(Optional) **firewall transparent** **Example:** `hostname(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Configuring the Firewall Mode" section on page 6-1 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode:	

	Routed mode:
	interface vlan *number*
	ip address *ip_address* [*mask*]
	nameif *name*
	security-level *level*
	
	Example:
	`hostname(config)# interface vlan 1`
	`hostname(config-if)# ip address 192.168.1.1 255.255.255.0`
	`hostname(config-if)# nameif inside`
	`hostname(config-if)# security-level 100`
		Configures an interface in routed mode. The **security-level** is a number between 1 and 100, where 100 is the most secure.

	Command	Purpose

	Transparent mode:
	interface bvi *number*
	  ip address *ip_address* [*mask*]
	
	interface vlan *number*
	  bridge-group *bvi_number*
	  nameif *name*
	  security-level *level*
	
	Example:
	`hostname(config)# interface bvi 1`
	`hostname(config-if)# ip address 192.168.1.1 255.255.255.0`
	
	`hostname(config)# interface vlan 1`
	`hostname(config-if)# bridge-group 1`
	`hostname(config-if)# nameif inside`
	`hostname(config-if)# security-level 100`
		Configures a bridge virtual interface and assigns a management VLAN to the bridge group. The **security-level** is a number between 1 and 100, where 100 is the most secure.

Step 3	(For directly-connected management hosts)
	dhcpd address *ip_address-ip_address interface_name*
	dhcpd enable *interface_name*
	
	Example:
	`hostname(config)# dhcpd address 192.168.1.2-192.168.1.254 inside`
	`hostname(config)# dhcpd enable inside`
		Enables DHCP for the management host on the management interface network. Make sure you do not include the management address in the range.

Step 4	(For remote management hosts) **route** *management_ifc management_host_ip mask gateway_ip* **1** **Example:** `hostname(config)# route management 10.1.1.0 255.255.255.0 192.168.1.50`	Configures a route to the management hosts.
Step 5	**http server enable** **Example:** `hostname(config)# http server enable`	Enables the HTTP server for ASDM.
Step 6	**http** *ip_address mask interface_name* **Example:** `hostname(config)# http 192.168.1.0 255.255.255.0 management`	Allows the management host to access ASDM.
Step 7	**write memory** **Example:** `hostname(config)# write memory`	Saves the configuration.

	Command	Purpose
Step 8	(Optional) **mode multiple** **Example:** `hostname(config)# mode multiple`	Sets the mode to multiple mode. When prompted, confirm that you want to convert the existing configuration to be the admin context. You are then prompted to reload the ASASM. See Chapter 9, "Configuring Multiple Context Mode," for more information.
Step 9	To launch ASDM, see the "Starting ASDM" section on page 3-13.	

Examples

The following routed mode configuration configures the VLAN 1 interface and enables ASDM for a management host:

```
interface vlan 1
    nameif inside
    ip address 192.168.1.1 255.255.255.0
    security-level 100
dhcpd address 192.168.1.3-192.168.1.254 inside
dhcpd enable inside
http server enable
http 192.168.1.0 255.255.255.0 inside
```

The following configuration converts the firewall mode to transparent mode, configures the VLAN 1 interface and assigns it to BVI 1, and enables ASDM for a management host:

```
firewall transparent
interface bvi 1
    ip address 192.168.1.1 255.255.255.0
interface vlan 1
    bridge-group 1
    nameif inside
    security-level 100
dhcpd address 192.168.1.3-192.168.1.254 inside
dhcpd enable inside
http server enable
http 192.168.1.0 255.255.255.0 inside
```

Starting ASDM

You can start ASDM using two methods:

- ASDM-IDM Launcher—The Launcher is an application downloaded from the ASA using a web browser that you can use to connect to any ASA IP address. You do not need to re-download the launcher if you want to connect to other ASAs. The Launcher also lets you run a virtual ASDM in Demo mode using files downloaded locally.
- Java Web Start—For each ASA that you manage, you need to connect with a web browser and then save or launch the Java Web Start application. You can optionally save the application to your PC; however you need separate applications for each ASA IP address.

Note Within ASDM, you can choose a different ASA IP address to manage; the difference between the Launcher and Java Web Start application functionality rests primarily in how you initially connect to the ASA and launch ASDM.

This section describes how to connect to ASDM initially, and then launch ASDM using the Launcher or the Java Web Start application. This section includes the following topics:

- Connecting to ASDM for the First Time, page 3-14
- Starting ASDM from the ASDM-IDM Launcher, page 3-15
- Starting ASDM from the Java Web Start Application, page 3-15
- Using ASDM in Demo Mode, page 3-16

Note ASDM allows multiple PCs or workstations to each have one browser session open with the same ASA software. A single ASA 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 ASA. In multiple context mode, five concurrent ASDM sessions are supported per context, up to a maximum of 32 total connections for each ASA.

Connecting to ASDM for the First Time

To connect to ASDM for the first time to download the ASDM-IDM Launcher or Java Web Start application, perform the following steps:

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

https://*interface_ip_address***/admin**

Where *interface_ip_address* is the management IP address of the ASA. See the "Configuring ASDM Access for Appliances" section on page 3-6 or the "Configuring ASDM Access for the ASA Services Module" section on page 3-11 for more information about management access.

See the ASDM release notes for your release for the requirements to run ASDM.

The ASDM launch page appears with the following buttons:

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

Step 2 To download the Launcher:

a. Click **Install ASDM Launcher and Run ASDM**.

b. Enter the username and password, and click **OK**. For a factory default configuration, leave these fields empty. With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. With HTTPS authentication enabled, enter your username and associated password.

c. Save the installer to your PC, and then start the installer. The ASDM-IDM Launcher opens automatically after installation is complete.

d. See the "Starting ASDM from the ASDM-IDM Launcher" section on page 3-15 to use the Launcher to connect to ASDM.

Step 3 To use the Java Web Start application:

a. Click **Run ASDM** or **Run Startup Wizard**.

b. Save the application to your PC when prompted. You can optionally open it instead of saving it.

c. See the "Starting ASDM from the Java Web Start Application" section on page 3-15 to use the Java Web Start application to connect to ASDM.

Starting ASDM from the ASDM-IDM Launcher

To start ASDM from the ASDM-IDM Launcher, perform the following steps.

Prerequisites

Download the ASDM-IDM Launcher according to the "Connecting to ASDM for the First Time" section on page 3-14.

Detailed Steps

Step 1 Start the ASDM-IDM Launcher application.

Step 2 Enter or choose the ASA 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**.

For a factory default configuration, leave these fields empty. With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. With HTTPS authentication enabled, enter your username and associated password.

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

The main ASDM window appears.

Starting ASDM from the Java Web Start Application

To start ASDM from the Java Web Start application, perform the following steps.

Prerequisites

Download the Java Web Start application according to the "Connecting to ASDM for the First Time" section on page 3-14.

Detailed Steps

Step 1 Start the Java Web Start application.

Step 2 Accept any certificates according to the dialog boxes that appear. The Cisco ASDM-IDM Launcher appears.

Step 3 Enter the username and password, and click **OK**. For a factory default configuration, leave these fields empty. With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. With HTTPS authentication enabled, enter your username and associated password.

The main ASDM window appears.

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 ASA features using the ASDM interface.
- Perform configuration and monitoring tasks with the CSC SSM.
- Obtain simulated monitoring and logging data, including real-time syslog 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 has been updated to support the following features:

- For global policies, an ASA in single, routed mode and intrusion prevention
- For object NAT, an ASA in single, routed mode and a firewall DMZ.
- For the Botnet Traffic Filter, an ASA in single, routed mode and security contexts.
- Site-to-Site VPN with IPv6 (Clientless SSL VPN and IPsec VPN)
- Promiscuous IDS (intrusion prevention)
- Unified Communication Wizard

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 the following location: http://www.cisco.com/cisco/web/download/index.html.

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.

Factory Default Configurations

The factory default configuration is the configuration applied by Cisco to new ASAs.

- ASA 5505—The factory default configuration configures interfaces and NAT so that the ASA is ready to use in your network immediately.
- ASA 5510 and higher—The factory default configuration configures an interface for management so you can connect to it using ASDM, with which you can then complete your configuration.

The factory default configuration is available only for routed firewall mode and single context mode. See Chapter 9, “Configuring Multiple Context Mode,” for more information about multiple context mode. See Chapter 6, “Configuring the Transparent or Routed Firewall,” for more information about routed and transparent firewall mode. For the ASA 5505, a sample transparent mode configuration is provided in this section.

Note In addition to the image files and the (hidden) default configuration, the following folders and files are standard in flash memory: log/, crypto_archive/, and coredumpinfo/coredump.cfg. The date on these files may not match the date of the image files in flash memory. These files aid in potential troubleshooting; they do not indicate that a failure has occurred.

This section includes the following topics:

- Restoring the Factory Default Configuration, page 3-18
- ASA 5505 Default Configuration, page 3-20
- ASA 5510 and Higher Default Configuration, page 3-24

Restoring the Factory Default Configuration

This section describes how to restore the factory default configuration.

Note On the ASASM, restoring the factory default configuration simply erases the configuration; there is no factory default configuration.

Limitations

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 ASA with a cleared configuration does not have any defined contexts to configure automatically using this feature.

Detailed Steps

Using the CLI:

	Command	Purpose
Step 1	**configure factory-default** [*ip_address* [*mask*]] **Example:** hostname(config)# **configure factory-default 10.1.1.1 255.255.255.0**	Restores the factory default configuration. If you specify the *ip_address*, then you set the inside or management interface IP address, depending on your model, instead of using the default IP address of 192.168.1.1. The **http** command uses the subnet you specify. Similarly, the **dhcpd address** command range consists of addresses within the subnet that you specify. **Note** This command also clears the **boot system** command, if present, along with the rest of the configuration. The **boot system** command lets you boot from a specific image, including an image on the external flash memory card. The next time you reload the ASA after restoring the factory configuration, it boots from the first image in internal flash memory; if you do not have an image in internal flash memory, the ASA does not boot.
Step 2	**write memory** **Example:** active(config)# write memory	Saves the default configuration to flash memory. This command saves the running configuration to the default location for the startup configuration, even if you previously configured the **boot config** command to set a different location; when the configuration was cleared, this path was also cleared.

Using ASDM:

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 (Optional) Enter the Management IP address of the management interface, instead of using the default address, 192.168.1.1. (For an ASA with a dedicated management interface, the interface is called "Management0/0.")

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

Step 4 Click **OK**.

A confirmation dialog box appears.

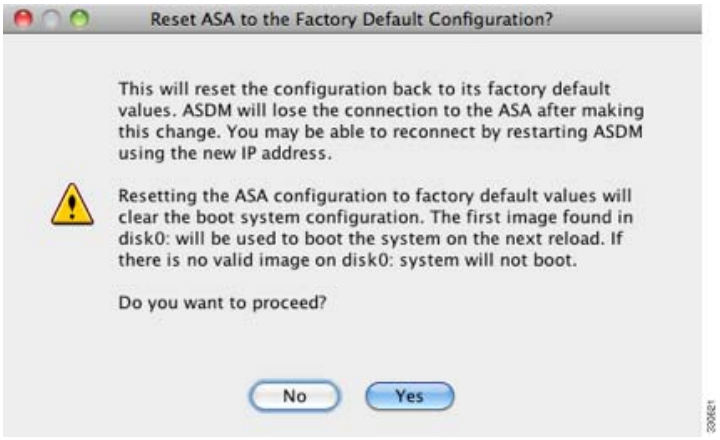

Note This action also clears the location of the boot image location, if present, along with the rest of the configuration. The Configuration > Device Management > System Image/Configuration > Boot Image/Configuration pane lets you boot from a specific image, including an image on the external memory. The next time you reload the ASA after restoring the factory configuration, it boots from the first image in internal flash memory; if you do not have an image in *internal* flash memory, the ASA does not boot.

Step 5 Click **Yes**.

Step 6 After you restore the default configuration, save this configuration to internal flash memory. Choose **File > Save Running Configuration to Flash**.

Choosing this option saves the running configuration to the default location for the startup configuration, even if you have previously configured a different location. When the configuration was cleared, this path was also cleared.

What to Do Next

See the "Getting Started with the Configuration" section on page 3-24 to start configuring the ASA.

ASA 5505 Default Configuration

The default configuration is available for routed mode only. This section describes the default configuration and also provides a sample transparent mode configuration that you can copy and paste as a starting point. This section includes the following topics:

- ASA 5505 Routed Mode Default Configuration, page 3-21
- ASA 5505 Transparent Mode Sample Configuration, page 3-22

ASA 5505 Routed Mode Default Configuration

The default factory configuration for the ASA 5505 configures the following:

- Interfaces—Inside (VLAN 1) and outside (VLAN 2).
- Switchports enabled and assigned—Ethernet 0/1 through 0/7 switch ports assigned to inside. Ethernet 0/0 assigned to outside.
- IP addresses— Outside address from DHCP; inside address set manually to 192.168.1.1/24.
- Network address translation (NAT)—All inside IP addresses are translated when accessing the outside using interface PAT.
- Traffic flow—IPv4 and IPv6 traffic allowed from inside to outside (this behavior is implicit on the ASA). Outside users are prevented from accessing the inside.
- DHCP server—Enabled for inside hosts, so a PC connecting to the inside interface receives an address between 192.168.1.5 and 192.168.1.254. DNS, WINS, and domain information obtained from the DHCP client on the outside interface is passed to the DHCP clients on the inside interface.
- Default route—Derived from DHCP.
- ASDM access—Inside hosts allowed.

Figure 3-1 ASA 5505 Routed Mode

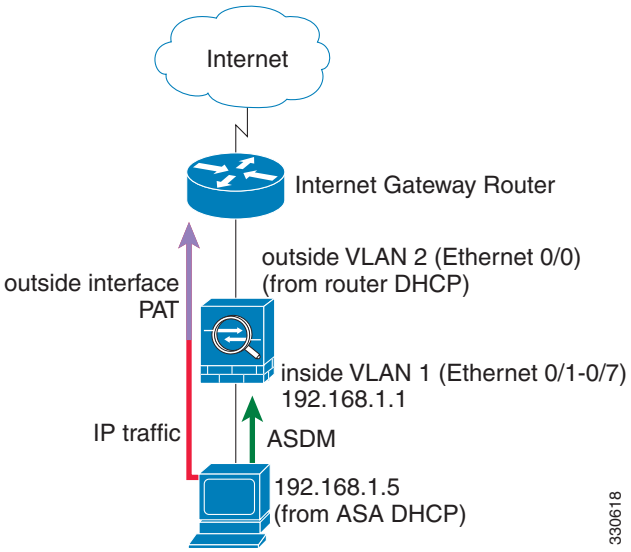

The configuration consists of the following commands:

```
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
object network obj_any
   subnet 0 0
   nat (inside,outside) dynamic interface
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.5-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
```

Note For testing purposes, you can allow ping from inside to outside by enabling ICMP inspection. Add the following commands to the default configuration:

```
policy-map global_policy
   class inspection_default
      inspect icmp
```

ASA 5505 Transparent Mode Sample Configuration

When you change the mode to transparent mode, the configuration is erased. You can copy and paste the following sample configuration at the CLI to get started. This configuration uses the default configuration as a starting point. Note the following areas you may need to modify:

- IP addresses—The IP addresses configured should be changed to match the network to which you are connecting.
- Static routes—For some kinds of traffic, static routes are required. See the "MAC Address vs. Route Lookups" section on page 6-4.

Figure 3-2 ASA 5505 Transparent Mode

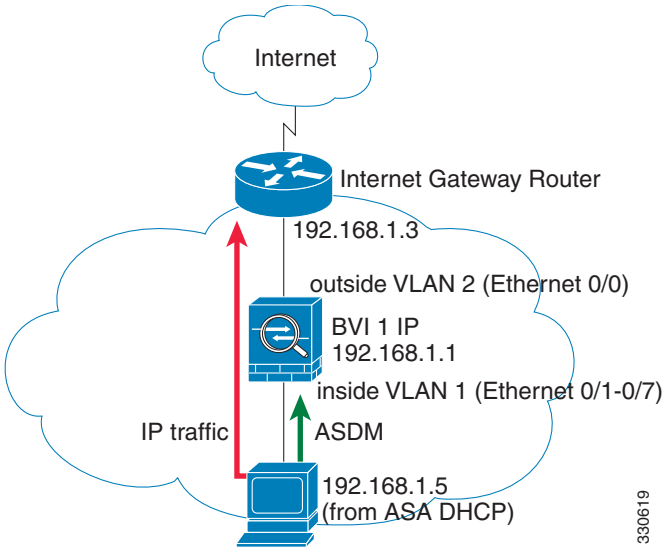


```
firewall transparent
interface Ethernet 0/0
   switchport access vlan 2
   no shutdown
interface Ethernet 0/1
   switchport access vlan 1
   no shutdown
interface Ethernet 0/2
   switchport access vlan 1
   no shutdown
interface Ethernet 0/3
   switchport access vlan 1
   no shutdown
interface Ethernet 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 bvi 1
   ip address 192.168.1.1 255.255.255.0
interface vlan2
   nameif outside
   security-level 0
   bridge-group 1
   no shutdown
interface vlan1
   nameif inside
   security-level 100
   bridge-group 1
   no shutdown
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.5-192.168.1.254 inside
```

```
dhcpd enable inside
```

Note For testing purposes, you can allow ping from inside to outside by enabling ICMP inspection. Add the following commands to the sample configuration:

```
policy-map global_policy
   class inspection_default
      inspect icmp
```

ASA 5510 and Higher Default Configuration

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

- Management interface—Management 0/0 (management).
- IP address—The management address is 192.168.1.1/24.
- DHCP server—Enabled for management hosts, so a PC connecting to the management interface receives an address between 192.168.1.2 and 192.168.1.254.
- ASDM access—Management hosts allowed.

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
```

Getting Started with the Configuration

To configure and monitor the ASA, 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**.

Step 5 To use 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 feature buttons to display the associated configuration pane.

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

Step 8 To monitor the ASA, click the **Monitoring** button on the toolbar and then click a feature button to display the associated monitoring pane.

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

Using the Command Line Interface Tool in ASDM

This section tells how to enter commands using ASDM, and how to work with the CLI. This section includes the following topics:

- Using the Command Line Interface Tool, page 3-25
- Handling Command Errors, page 3-26
- Using Interactive Commands, page 3-26
- Avoiding Conflicts with Other Administrators, page 3-26
- Showing Commands Ignored by ASDM on the Device, page 3-26

Using the Command Line Interface Tool

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

The commands you can enter with the CLI tool depend on your user privileges. See the "Information About Authorization" section on page 46-2 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 ASA.

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.

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 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 ASA. Before using the ASDM CLI tool to make configuration changes, check for other active administrative sessions. If more than one user is configuring the ASA at the same time, the most recent changes take effect.

To view other administrative sessions that are currently active on the same ASA, 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 4-34 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.

Applying Configuration Changes to Connections

When you make security policy changes to the configuration, all *new* connections use the new security policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. **show** command output for old connections reflect the old configuration, and in some cases will not include data about the old connections.

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

To ensure that all connections use the new policy, you need to disconnect the current connections so they can reconnect using the new policy.

To disconnect connections, enter one of the following commands.

Detailed Steps

Command	Purpose
clear local-host [*ip_address*] [**all**] **Example:** `hostname(config)# clear local-host all`	This command reinitializes per-client run-time states such as connection limits and embryonic limits. As a result, this command removes any connection that uses those limits. See the **show local-host all** command to view all current connections per host. With no arguments, this command clears all affected through-the-box connections. To also clear to-the-box connections (including your current management session), use the **all** keyword. To clear connections to and from a particular IP address, use the *ip_address* argument.
clear conn [**all**] [**protocol** {**tcp** | **udp**}] [**address** *src_ip*[**-***src_ip*] [**netmask** *mask*]] [**port** *src_port*[**-***src_port*]] [**address** *dest_ip*[**-***dest_ip*] [**netmask** *mask*]] [**port** *dest_port*[**-***dest_port*]] **Example:** `hostname(config)# clear conn all`	This command terminates connections in any state. See the **show conn** command to view all current connections. With no arguments, this command clears all through-the-box connections. To also clear to-the-box connections (including your current management session), use the **all** keyword. To clear specific connections based on the source IP address, destination IP address, port, and/or protocol, you can specify the desired options.

Using ASDM

This chapter describes how to use the ASDM user interface, and includes the following sections:

- Information About the ASDM User Interface, page 4-1
- Navigating in the ASDM User Interface, page 4-3
- Menus, page 4-4
- Toolbar, page 4-10
- ASDM Assistant, page 4-10
- Status Bar, page 4-11
- Device List, page 4-11
- Common Buttons, page 4-12
- Keyboard Shortcuts, page 4-12
- Find Function, page 4-14
- Enabling Extended Screen Reader Support, page 4-16
- Organizational Folder, page 4-16
- About the Help Window, page 4-16
- Home Pane (Single Mode and Context), page 4-17
- Home Pane (System), page 4-31
- Defining ASDM Preferences, page 4-32
- Using the ASDM Assistant, page 4-33
- Enabling History Metrics, page 4-34
- Unsupported Commands, page 4-34

Information About the ASDM User Interface

The ASDM user interface is designed to provide easy access to the many features that the ASA supports. The ASDM user interface includes the following elements:

- A menu bar that provides quick access to files, tools, wizards, and help. Many menu items also have keyboard shortcuts.
- A toolbar that enables you to navigate ASDM. From the toolbar you can access the home, configuration, and monitoring panes. You can also get help and navigate between panes.

- A dockable left Navigation pane to move through the Configuration and Monitoring panes. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it. To access the Configuration and Monitoring panes, you can do one of the following:
  - Click links on the left side of the application window in the left Navigation pane. The Content pane then displays the path (for example, Configuration > Device Setup > Startup Wizard) in the title bar of the selected pane.
  - If you know the exact path, you can type it directly into the title bar of the Content pane on the right side of the application window, without clicking any links in the left Navigation pane.
- A maximize and restore button in the right corner of the Content pane that lets you hide and show the left Navigation pane.
- A dockable device list pane with a list of devices that you can access through ASDM. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it. For more information, see the "Device List" section on page 4-11.
- A status bar that shows the time, connection status, user, memory status, running configuration status, privilege level, and SSL status at the bottom of the application window.
- A left Navigation pane that shows various objects that you can use in the rules tables when you create access rules, NAT rules, AAA rules, filter rules, and service rules. The tab titles within the pane change according to the feature that you are viewing. In addition, the ASDM Assistant appears in this pane.

Figure 4-1 on page 4-2 shows the elements of the ASDM user interface.

Figure 4-1 ASDM User Interface

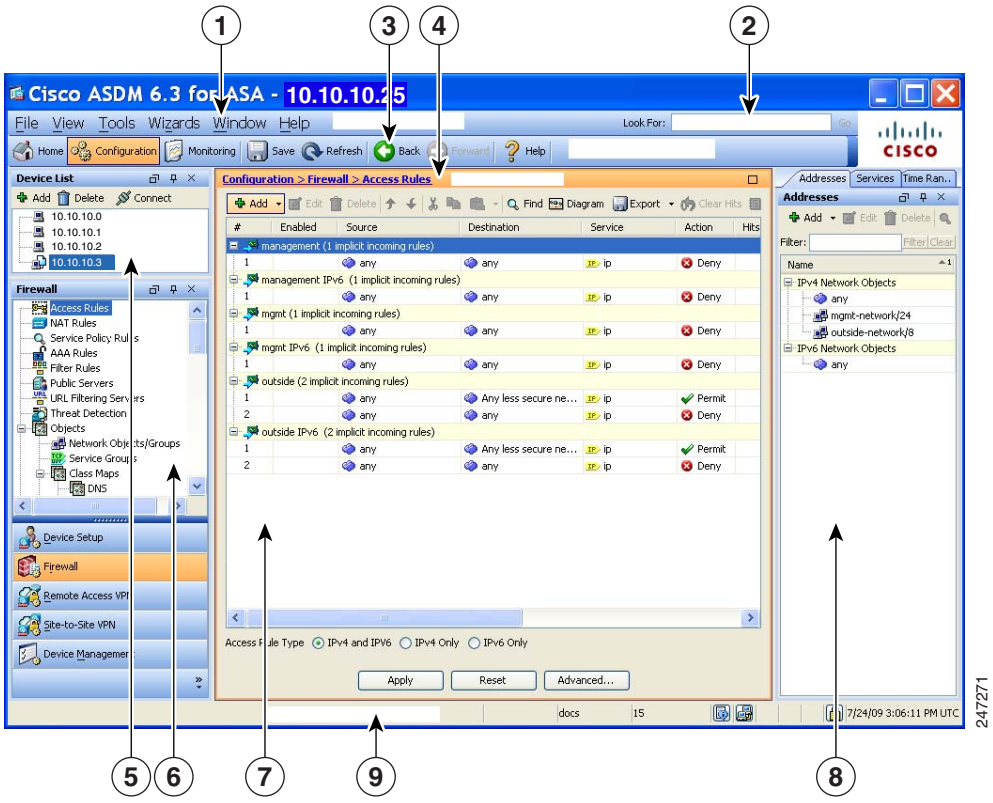

Legend

GUI Element	Description
1	Menu Bar
2	Search Field
3	Toolbar
4	Navigation Path
5	Device List Pane
6	Left Navigation Pane
7	Content Pane
8	Right Navigation Pane
9	Status Bar

Note Tool tips have been added for various parts of the GUI, including Wizards, the Configuration and Monitoring panes, and the Status Bar. To view tool tips, hover your mouse over a specific user interface element, such as an icon in the status bar.

Navigating in the ASDM User Interface

To move efficiently throughout the ASDM user interface, you may use a combination of menus, the toolbar, dockable panes, and the left and right Navigation panes, which are described in the previous section. The available functions appear in a list of buttons below the Device List pane. An example list could include the following function buttons:

- Device Setup
- Firewall
- Trend Micro Content Security
- Botnet Traffic Filter
- Remote Access VPN
- Site to Site VPN
- Device Management

The list of function buttons that appears is based on the licensed features that you have purchased. Click each button to access the first pane in the selected function for either the Configuration view or the Monitoring view. The function buttons are not available in the Home view.

To change the display of function buttons, perform the following steps:

Step 1 Choose the drop-down list below the last function button to display a context menu.

Step 2 Choose one of the following options:

- To show more buttons, click **Show More Buttons**.
- To show fewer buttons, click **Show Fewer Buttons**.

- To add or remove buttons, click **Add or Remove Buttons**, then click the button to add or remove from the list that appears.
- To change the sequence of the buttons, choose **Option** to display the Option dialog box, which displays a list of the buttons in their current order. Then choose one of the following:
  - To move up a button in the list, click **Move Up**.
  - To move down a button in the list, click **Move Down**.
  - To return the order of the items in the list to the default setting, click **Reset**.

Step 3 To save your settings and close this dialog box, click **OK**.

Menus

You can access ASDM menus using the mouse or keyboard. For information about accessing the menu bar from the keyboard, see the "Keyboard Shortcuts" section on page 4-12.

ASDM has the following menus:

- File Menu, page 4-4
- View Menu, page 4-5
- Tools Menu, page 4-6
- Wizards Menu, page 4-8
- Window Menu, page 4-9
- Help Menu

File Menu

The File menu lets you manage ASA configurations. The following table lists the tasks that you can perform using the File menu.

File Menu Item	Description
Refresh ASDM with the Running Configuration on the Device	Loads a copy of the running configuration into ASDM.
Refresh	Ensures that ASDM has a current copy of the running configuration.
Reset Device to the Factory Default Configuration	Restores the configuration to the factory default. See the "Restoring the Factory Default Configuration" section on page 3-18 for more information.
Show Running Configuration in New Window	Displays the current running configuration in a new window.
Save Running Configuration to Flash	Writes a copy of the running configuration to flash memory.

File Menu Item	Description
Save Running Configuration to TFTP Server	Stores a copy of the current running configuration file on a TFTP server. See the "Saving the Running Configuration to a TFTP Server" section on page 101-27 for more information.
Save Running Configuration to Standby Unit	Sends a copy of the running configuration file on the primary unit to the running configuration of a failover standby unit.
Save Internal Log Buffer to Flash	Saves the internal log buffer to flash memory.
Print	Prints the current page. We recommend landscape page orientation when you print rules. When you use Internet Explorer, permission to print was already granted when you originally accepted the signed applet.
Clear ASDM Cache	Removes local ASDM images. ASDM downloads images locally when you connect to ASDM.
Clear ASDM Password Cache	Removes the password cache if you have defined a new password and still have a existing password that is different than the new password.
Clear Internal Log Buffer	Empties the syslog message buffer.
Exit	Closes ASDM.

View Menu

The View menu lets you display various parts of the ASDM user interface. Certain items are dependent on the current view. You cannot select items that cannot be displayed in the current view. The following table lists the tasks that you can perform using the View menu.

View Menu Item	Description
Home	Displays the Home view.
Configuration	Displays the Configuration view.
Monitoring	Displays the Monitoring view.
Device List	Display a list of devices in a dockable pane. See the "Device List" section on page 4-11 for more information.
Navigation	Shows and hides the display of the Navigation pane in the Configuration and Monitoring views.
ASDM Assistant	Searches and finds useful ASDM procedural help about certain tasks. See the "ASDM Assistant" section on page 4-10 for more information.
SIP Details	Shows and hides voice network information.
Latest ASDM Syslog Messages	Shows and hides the display of the Latest ASDM Syslog Messages pane in the Home view. This pane is only available in the Home view. If you do not have sufficient memory to upgrade to the most current release, syslog message %ASA-1-211004 is generated, indicating what the installed memory is and what the required memory is. This message reappears every 24 hours until the memory is upgraded.

View Menu Item	Description
Addresses	Shows and hides the display of the Addresses pane. The Addresses pane is only available for the Access Rules, NAT Rules, Service Policy Rules, AAA Rules, and Filter Rules panes in the Configuration view.
Services	Shows and hides the display of the Services pane. The Services pane is only available for the Access Rules, NAT Rules, Service Policy Rules, AAA Rules, and Filter Rules panes in the Configuration view.
Time Ranges	Shows and hides the display of the Time Ranges pane. The Time Ranges pane is only available for the Access Rules, Service Policy Rules, AAA Rules, and Filter Rules panes in the Configuration view.
Global Pools	Shows and hides the display of the Global Pools pane. The Global Pools pane is only available for the NAT Rules pane in the Configuration view.
Find in ASDM	Locates an item for which you are searching, such as a feature or the ASDM Assistant.
Back	Returns to the previous pane. See the "Common Buttons" section on page 4-12 for more information.
Forward	Goes to the next pane previously visited. See the "Common Buttons" section on page 4-12 for more information.
Reset Layout	Returns the layout to the default configuration.
Office Look and Feel	Changes the screen fonts and colors to the Microsoft Office settings.

Tools Menu

The Tools menu provides you with the following series of tools to use in ASDM.

Tools Menu Item	Description
Command Line Interface	Sends commands to the ASA and view the results. See the "Using the Command Line Interface Tool in ASDM" section on page 3-25 for more information.
Show Commands Ignored by ASDM on Device	Displays unsupported commands that have been ignored by ASDM. See the "Showing Commands Ignored by ASDM on the Device" section on page 3-26 for more information.
Packet Tracer	Traces a packet from a specified source address and interface to a destination. You can specify the protocol and port of any type of data and view the lifespan of a packet, with detailed information about actions taken on it. See the "Tracing Packets with Packet Tracer" section on page 73-7 for more information.

Tools Menu Item	Description
Ping	Verifies the configuration and operation of the ASA and surrounding communications links, as well as performs basic testing of other network devices. See the "Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping" section on page 73-3 for more information.
Traceroute	Determines the route that packets will take to their destination. See the "Determining Packet Routing with Traceroute" section on page 73-6 for more information.
File Management	Views, moves, copies, and deletes files stored in flash memory. You can also create a directory in flash memory. See the "Managing Files" section on page 101-14 for more information. You can also transfer files between various file systems, including TFTP, flash memory, and your local PC. See the "Transferring Files" section on page 101-17 for more information.
Upgrade Software from Local Computer	Uploads an ASA image, ASDM image, or another image on your PC to flash memory. See the "Managing Files" section on page 101-14 dialog box for more information.
Check for ASA/ASDM Updates	Upgrades ASA software and ASDM software through a wizard. See the "The Upgrade Software from Cisco.com Wizard lets you automatically upgrade the ASDM and ASA to more current versions." section on page 101-4 for more information.
Backup Configurations	Backs up the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles. See the "Backing Up Configurations" section on page 101-21 for more information.
Restore Configurations	Restores the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles. See the "Restoring Configurations" section on page 101-25 for more information.
System Reload	Restarts the ASDM and reload the saved configuration into memory. See the "Scheduling a System Restart" section on page 101-27 for more information.
Administrator's Alerts to Clientless SSL VPN Users	Enables an administrator to send an alert message to clientless SSL VPN users. See the "Sending an Administrator's Alert to Clientless SSL VPN Users" section on page 91-149 for more information.
Preferences	Changes the behavior of specified ASDM functions between sessions. See the "Defining ASDM Preferences" section on page 4-32 for more information.
ASDM Java Console	Shows the Java console. See the "Viewing and Copying Logged Entries with the ASDM Java Console" section on page 96-13 for more information.

Wizards Menu

The Wizards menu lets you run a wizard to configure multiple features. The following table lists the available Wizards and their features.

Wizards Menu Item	Description
Startup Wizard	Guides you, step-by-step, through the initial configuration of the ASA. For more information, see Chapter 7, "Using the Startup Wizard.".
IPsec VPN Wizard	Enables you to configure an IPsec VPN policy on the ASA. For more information, see Chapter 85, "VPN Wizards.".
SSL VPN Wizard	Enables you to configure an SSL VPN policy on the ASA. For more information, see Chapter 85, "VPN Wizards.".
High Availability and Scalability Wizard	Allows you to configure failover: VPN cluster load balancing, or ASA clustering on the ASA. For more information, see: • Active/Standby failover—See the "Using the High Availability and Scalability Wizard" section on page 12-7. • Active/Active failover—See the "Task Flow for Using the High Availability and Scalability Wizard" section on page 13-8. • VPN cluster load balancing—See the "Configuring VPN Cluster Load Balancing with the High Availability and Scalability Wizard" section on page 86-25. • ASA clustering—See the "Adding or Joining an ASA Cluster" section on page 10-44.
Unified Communication Wizard	Enables you to configure unified communication features, such as an IP phone, on the ASA. For more information, see Chapter 64, "Information about the Cisco Unified Communication Wizard.".
Packet Capture Wizard	Allows you to configure packet capture on the ASA. The wizard runs one packet capture on each ingress and egress interface. After you run the capture, you can save it on your computer, and then examine and analyze the capture with a packet analyzer. For more information, see the "Configuring and Running Captures with the Packet Capture Wizard" section on page 102-1.

Window Menu

The Window menu enables you to move between ASDM windows. The active window appears as the selected window.

Help Menu

The Help menu provides links to online Help, as well as information about ASDM and the ASA. The following table lists the tasks that you can perform using the Help menu.

Help Menu Items	Description
Help Topics	Opens a new browser window with help organized by contents, window name, and indexed in the left frame. Use these methods to find help for any topic, or search using the Search tab.
Help for Current Screen	Opens context-sensitive help about that screen. The wizard runs the screen, pane, or dialog box that is currently open. Alternatively, you can also click the question mark (**?**) help icon.
Release Notes	Opens the most current version of the *ASDM release notes* on Cisco.com. The release notes contain the most current information about ASDM software and hardware requirements, and the most current information about changes in the software.
ASDM Assistant	Opens the ASDM Assistant, which lets you search downloadable content from Cisco.com, with details about performing certain tasks.
About Cisco Adaptive Security Appliance (ASA)	Displays information about the ASA, including the software version, hardware set, configuration file loaded at startup, and software image loaded at startup. This information is helpful in troubleshooting.
About Cisco ASDM 6.3	Displays information about ASDM such as the software version, hostname, privilege level, operating system, device type, and Java version.

Toolbar

The Toolbar below the menus provides access to the Home view, Configuration view, and Monitoring view. It also lets you choose between the system and security contexts in multiple context mode, and provides navigation and other commonly used features. The following table lists the tasks that you can perform using the Toolbar.

Toolbar Button	Description
System/Contexts	Shows which context you are in. To open the context list in the left-hand pane, click the down arrow, then click the up arrow to restore the context drop-down list. After you have expanded this list, click the left arrow to collapse the pane, then the right arrow to restore the pane. To manage the system, choose **System** from the drop-down list. To manage the context, choose one from the drop-down list.
Home	Displays the Home pane, which lets you view important information about your ASA such as the status of your interfaces, the version you are running, licensing information, and performance. See the "Home Pane (Single Mode and Context)" section on page 4-17 for more information. In multiple mode, the system does not have a Home pane.
Configuration	Configures the ASA. Click a function button in the left Navigation pane to configure that function.
Monitoring	Monitors the ASA. Click a function button in the left Navigation pane to configure that function.
Back	Returns to the last pane of ASDM that you visited.
Forward	Goes forward to the last pane of ASDM that you visited.
Search	Searches for a feature in ASDM. The Search function looks through the titles of each pane and presents you with a list of matches, and gives you a hyperlink directly to that pane. If you need to switch quickly between two different panes that you found, click **Back** or **Forward**. See the "ASDM Assistant" section on page 4-10 for more information.
Refresh	Refreshes ASDM with the current running configuration, except for graphs in any of the Monitoring panes.
Save	Saves the running configuration to the startup configuration for write-accessible contexts only.
Help	Shows context-sensitive help for the screen that is currently open.

ASDM Assistant

The ASDM Assistant lets you search and view useful ASDM procedural help about certain tasks. This feature is available in routed and transparent modes, and in the single and system contexts.

To access information, choose **View > ASDM Assistant > How Do I?** or enter a search request from the Look For field in the menu bar. From the Find drop-down list, choose **How Do I?** to begin the search.

To use the ASDM Assistant, perform the following steps:

Step 1 In the main ASDM application window, choose **View > ASDM Assistant**.

The ASDM Assistant pane appears.

Step 2 In the Search field, enter the information that you want to find, and click **Go**.

The requested information appears in the Search Results pane.

Step 3 Click any links that appear in the Search Results and Features areas to obtain more details.

Status Bar

The status bar appears at the bottom of the ASDM window. The following table lists the areas shown from left to right.

Area	Description
Status	The status of the configuration (for example, "Device configuration loaded successfully.")
Failover	The status of the failover unit, either active or standby.
User Name	The username of the ASDM user. If you logged in without a username, the username is "admin."
User Privilege	The privilege of the ASDM user.
Commands Ignored by ASDM	Click the icon to show a list of commands from your configuration that ASDM did not process. These commands will not be removed from the configuration.
Connection to Device	The ASDM connection status to the ASA. See the "Connection to Device" section on page 4-11 for more information.
Syslog Connection	The syslog connection is up, and the ASA is being monitored.
SSL Secure	The connection to ASDM is secure because it uses SSL.
Time	The time that is set on the ASA.

Connection to Device

ASDM maintains a constant connection to the ASA to maintain up-to-date Monitoring and Home pane data. This dialog box shows the status of the connection. When you make a configuration change, ASDM opens a second connection for the duration of the configuration, and then closes it; however, this dialog box does not represent the second connection.

Device List

The device list is a dockable pane. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it. This pane is available in the Home, Configuration, Monitoring, and System views. You can use this pane to switch to another device; however, that device must run the same version of ASDM that you are currently running. To display the pane fully, you must have at least two devices listed. This feature is available in routed and transparent modes, and in the single, multiple, and system contexts.

To use this pane to connect to another device, perform the following steps:

Step 1 Click **Add** to add another device to the list.

The Add Device dialog box appears.

Step 2 In the Device/IP Address/Name field, type the device name or IP address of the device, and then click **OK**.

Step 3 Click **Delete** to remove a selected device from the list.

Step 4 Click **Connect** to connect to another device.

The Enter Network Password dialog box appears.

Step 5 Type your username and password in the applicable fields, and then click **Login**.

Common Buttons

Many ASDM panes include buttons that are listed in the following table. Click the applicable button to complete the desired task:

Button	Description
Apply	Sends changes made in ASDM to the ASA and applies them to the running configuration.
Save	Writes a copy of the running configuration to flash memory.
Reset	Discards changes and reverts to the information displayed before changes were made or the last time that you clicked Refresh or Apply. After you click **Reset**, click **Refresh** to make sure that information from the current running configuration appears.
Restore Default	Clears the selected settings and returns to the default settings.
Cancel	Discards changes and returns to the previous pane.
Enable	Displays read-only statistics for a feature.
Close	Closes an open dialog box.
Clear	Remove information from a field, or remove a check from a check box.
Back	Returns to the previous pane.
Forward	Goes to the next pane.
Help	Displays help for the selected pane or dialog box.

Keyboard Shortcuts

You can use the keyboard to navigate the ASDM user interface.

Table 4-1 lists the keyboard shortcuts you can use to move across the three main areas of the ASDM user interface.

Table 4-1 Keyboard Shortcuts Within the Main Window

To display the	Windows/Linux	MacOS
Home Pane	Ctrl+H	Shift+Command+H
Configuration Pane	Ctrl+G	Shift+Command+G
Monitoring Pane	Ctrl+M	Shift+Command+M
Help	F1	Command+?
Back	Alt+Left Arrow	Command+[
Forward	Alt+Rightarrow	Command+]
Refresh the display	F5	Command+R
Cut	Ctrl+X	Command+X
Copy	Ctrl+C	Command+C
Paste	Ctrl+V	Command+V
Save the configuration	Ctrl+S	Command+S
Popup menus	Shift+F10	—
Close a secondary window	Alt+F4	Command+W
Find	Ctrl+F	Command+F
Exit	Alt+F4	Command+Q
Exit a table or text area	Ctrl_Shift or Ctrl+Shift+Tab	Ctril+Shift or Ctrl+Shift+Tab

Table 4-2 lists the keyboard shortcut you can use to navigate within a pane.

Table 4-2 Keyboard Shortcuts Within a Pane

To move the focus to the	Press
Next field	Tab
Previous field	Shift+Tab
Next field when the focus is in a table	Ctrl+Tab
Previous field when the focus is in a table	Shift+Ctrl+Tab
Next tab (when a tab has the focus)	Right Arrow
Previous tab (when a tab has the focus)	Left Arrow
Next cell in a table	Tab
Previous sell in a table	Shift+Tab
Next pane (when multiple panes are displayed)	F6
Previous pane (when multiple panes are displayed)	Shift+F6

Table 4-3 lists the keyboard shortcuts you can use with the Log Viewers.

Table 4-3 Keyboard Shortcuts for the Log Viewer

To	Windows/Linux	MacOS
Pause and Resume Real-Time Log Viewer	Ctrl+U	Command+
Refresh Log Buffer Pane	F5	Command+R
Clear Internal Log Buffer	Ctrl+Delete	Command+Delete
Copy Selected Log Entry	Ctrl+C	Command+C
Save Log	Ctrl+S	Command+S
Print	Ctrl+P	Command+P
Close a secondary window	Alt+F4	Command+W

Table 4-4 lists the keyboard shortcuts you can use to access menu items.

Table 4-4 Keyboard Shortcuts to Access Menu Items

To access the	Windows/Linux
Menu Bar	Alt
Next Menu	Right Arrow
Previous Menu	Left Arrow
Next Menu Option	Down Arrow
Previous Menu Option	Up Arrow
Selected Menu Option	Enter

Find Function

This section includes the following topics:

- Using the Find Function in Most ASDM Panes, page 4-14d
- Using the Find Function in the ACL Manager Pane, page 4-15

Using the Find Function in Most ASDM Panes

Some ASDM panes contain tables with many elements. To make it easier for you to search, highlight, and then edit a particular entry, several ASDM panes have a find function that allows you to search on objects within those panes.

To perform a search, you can type a phrase into the Find field to search on all columns within any given pane. The phrase can contain the wild card characters “*” and “?”. The * matches one or more characters, and ? matches one character. The up and down arrows to the right of the Find field locate the next (up) or previous (down) occurrence of the phrase. Check the Match Case check box to find entries with the exact uppercase and lowercase characters that you enter.

For example, entering B*ton-L* might return the following matches:

```
Boston-LA, Boston-Lisbon, Boston-London
```

Entering Bo?ton might return the following matches:

```
Boston, Bolton
```

The following list shows the ASDM panes in which you can use the find function:

- AAA Server Groups panes
- ACL Manager panes—The find function in the ACL Manager pane differs from that of the other panes. See the "Using the Find Function in the ACL Manager Pane" section on page 4-15 for more information.
- Certificate-to-Conn Profile Maps-Rules pane
- DAP panes
- Identity Certificates pane
- IKE Policies pane
- IPSec Proposals (Transform Sets) pane
- Local User panes
- Portal-Bookmark pane
- Portal-Customization panes
- Portal-Port Forwarding pane
- CA Certificates pane
- Portal-Smart Tunnels pane
- Portal-Web Contents pane
- VPN Connection Profiles panes
- VPN Group Policies panes

Using the Find Function in the ACL Manager Pane

Because ACLs and ACEs contain many elements of different types, the find function in the ACL Manager pane allows for a more targeted search than the find function in other panes.

To find elements within the ACL Manager pane, perform the following steps:

Step 1 In the ACL Manager pane, click **Find**.

Step 2 In the Filter field, choose one of the following options from the drop-down list:

- Source—The search includes a source IP address of a the network object group, interface IP, or any address from which traffic is permitted or denied. You specify this address in Step 4.
- Destination—The search includes a destination IP address (host or network) that is permitted or denied to send traffic to the IP addresses listed in the Source section. You specify this address in Step 4.
- Source or Destination—The search includes either a source or a destination address that you specify in Step 4.
- Service—The search includes a service group or predefined service policy that you specify in Step 4.

- Query—When you choose Query from the drop-down list, click **Query** to specify a detailed search by all four of the preceding options: Source, Destination, Source or Destination, and Service.

Step 3 In the second field, choose one of the following options from the drop-down list:

- is—Specifies an exact match of the detail that you enter in Step 4.
- contains—Specifies to search for ACLs or ACEs that contain, but are not limited to, the detail you enter in Step 4.

Step 4 In the third field, enter specific criteria about ACLs or ACEs that you would like to find, or click the browse button to search for key elements in your ACL/ACE configuration.

Step 5 Click **Filter** to perform the search.

The ASDM find function returns a list of ACLs and ACEs that contain your specified criteria.

Step 6 Click **Clear** to clear the list of found ACLs and ACEs.

Step 7 Click the red **x** to close the find function box.

Enabling Extended Screen Reader Support

By default, labels and descriptions are not included in tab order when you press the Tab key to navigate a pane. Some screen readers, such as JAWS, only read screen objects that have the focus. You can include the labels and descriptions in the tab order by enabling extended screen reader support.

To enable extended screen reader support, perform the following steps:

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

The Preferences dialog box appears.

Step 2 On the General tab, check the **Enable screen reader support** check box.

Step 3 Click **OK**.

Step 4 Restart ASDM to activate screen reader support.

Organizational Folder

Some folders in the navigation pane for the configuration and monitoring views do not have associated configuration or monitoring panes. These folders are used to organize related configuration and monitoring tasks. Clicking these folders displays a list of sub-items in the right Navigation pane. You can click the name of a sub-item to go to that item.

About the Help Window

This section includes the following topics:

- Header Buttons, page 4-17
- Browser Window, page 4-17

Header Buttons

To obtain the information that you need, click the applicable button listed in the following table.

Button	Description
About ASDM	Displays information about ASDM, including the hostname, version number, device type, ASA software version number, privilege level, username, and operating system being used.
Search	Searches for information among online help topics.
Using Help	Describes the most efficient methods for using online help.
Glossary	Lists terms found in ASDM and ASAs.
Contents	Displays a table of contents.
Screens	Lists help files by screen name.
Index	Displays an index of help topics found in ASDM online help.

Browser Window

When you open help and a help page is already open, the new help page will appear in the same browser window. If no help page is open, then the help page will appear in a new browser window.

When you open help and Netscape Communicator is the default browser, the help page will appear in a new browser window. If Internet Explorer is the default browser, the help page may appear either in the last-visited browser window or in a new browser window, according to your browser settings. You can control this behavior in Internet Explorer by choosing **Tools > Internet Options > Advanced > Reuse windows for launching shortcuts**.

Home Pane (Single Mode and Context)

The ASDM Home pane lets you view important information about your ASA. Status information in the home pane is updated every ten seconds. This pane usually has two tabs: Device Dashboard and Firewall Dashboard.

If you have a CSC SSM installed in your ASA, the Content Security tab also appears in the Home pane. The additional tab displays status information about the CSC SSM software.

If you have IPS software installed in your ASA, the Intrusion Prevention tab also appears in the Home pane. The additional tab displays status information about the IPS software.

This section includes the following topics:

- Device Dashboard Tab, page 4-18
- Firewall Dashboard Tab, page 4-22
- Cluster Dashboard Tab, page 4-25
- Cluster Firewall Dashboard Tab, page 4-26
- Content Security Tab, page 4-27
- Intrusion Prevention Tab, page 4-28
- ASA CX Status Tab, page 4-30

Device Dashboard Tab

The Device Dashboard tab lets you view, at a glance, important information about your ASA, such as the status of your interfaces, the version you are running, licensing information, and performance.

Figure 4-2 shows the elements of the Device Dashboard tab.

Figure 4-2 Device Dashboard Tab

Legend

GUI Element	Description
1	Device Information Pane, page 4-19
2	Interface Status Pane, page 4-20
3	VPN Sessions Pane, page 4-20
4	Traffic Status Pane, page 4-20
5	System Resources Status Pane, page 4-20
6	Traffic Status Pane, page 4-20
See Figure 4-3	Latest ASDM Syslog Messages Pane, page 4-20

Device Information Pane

The Device Information pane includes two tabs that show device information: General tab and License tab. Under the General tab you have access to the Environment Status button, which provides an at-a-glance view of the system health:

- General Tab, page 4-19
- License Tab, page 4-19

General Tab

This tab shows basic information about the ASA:

- Host name—Shows the hostname of the device.
- ASA version—Lists the version of ASA software that is running on the device.
- ASDM version—Lists the version of ASDM software that is running on the device.
- Firewall mode—Shows the firewall mode in which the device is running.
- Total flash—Displays the total RAM that is currently being used.
- ASA Cluster Role—When you enable clustering, shows the role of this unit, either Master or Slave.
- Device uptime—Shows the time in which the device has been operational since the latest software upload.
- Context mode—Shows the context mode in which the device is running.
- Total Memory—Shows the DRAM installed on the ASA.
- Environment status—Shows the system health. The ASA 5580 and 5585 chassis models provide a set of hardware statistics that is available by clicking the plus sign (+) to the right of the Environment Status label in the General tab. You can see how many power supplies are installed, track the operational status of the fan and power supply modules, and track the temperatures of the CPUs and the ambient temperature of the system.

  In general, the Environment Status button provides an at-a-glance view of the system health. If all monitored hardware components within the system are operating within normal ranges, the plus sign (+) button shows OK in green. Conversely, if any one component within the hardware system is operating outside of normal ranges, the plus sign (+) button turns into a red circle to show Critical status and to indicate that a hardware component requires immediate attention.

  For more information about specific hardware statistics, see the hardware guide for your particular device.

Note If you do not have enough memory to upgrade to the most current release of the ASA, the Memory Insufficient Warning dialog box appears. Follow the directions that appear in this dialog box to continue using the ASA and ASDM in a supported manner. Click **OK** to close this dialog box.

License Tab

This tab shows a subset of licensed features. To view detailed license information, or to enter a new activation key, click **More Licenses**; the Configuration > Device Management > Licensing > Activation Key pane appears. See Chapter 5, “Managing Feature Licenses.”

Interface Status Pane

This pane shows the status of each interface. If you select an interface row, the input and output throughput in Kbps displays below the table.

VPN Sessions Pane

This pane shows the VPN tunnel status. Click **Details** to go to the Monitoring > VPN > VPN Statistics > Sessions pane.

Failover Status Pane

This pane shows the failover status.

Click **Configure** to start the High Availability and Scalability Wizard. After you have completed the wizard, the failover configuration status (either Active/Active or Active/Standby) appears.

If failover is configured, click **Details** to open the Monitoring > Properties > Failover > Status pane.

System Resources Status Pane

This pane shows CPU and memory usage statistics.

Traffic Status Pane

This pane shows graphs for connections per second for all interfaces and for the traffic throughput of the lowest security interface.

When your configuration contains multiple lowest security level interfaces, and any one of them is named "outside," then that interface is used for the traffic throughput graphs. Otherwise, ASDM picks the first interface from the alphabetical list of lowest security level interfaces.

Latest ASDM Syslog Messages Pane

This pane shows the most recent system messages generated by the ASA, up to a maximum of 100 messages. If logging is disabled, click **Enable Logging** to enable logging.

Figure 4-3 shows the elements of the Latest ASDM Syslog Messages pane.

Figure 4-3 Latest ASDM Syslog Messages Pane

Legend

GUI Element	Description
1	To resize the pane, drag the divider up or down.
2	Expands the pane. To return the pane to the default size, click the double-square icon.
3	Makes a floating pane. To dock the pane, click the docked pane icon.
4	Enables or disables Auto-hide. When Auto-hide is enabled, move your cursor over the Latest ASDM Syslog Messages button in the left, bottom corner and the pane displays. Move your cursor away from the pane, and it disappears.
5	Closes the pane. To show the pane, choose **View Latest ASDM Syslog Messages.**
6	To continue updating the display of syslog messages, click the green icon on the right-hand side.
7	To stop updating the display of syslog messages, click the red icon on the right-hand side.
8	To open the Logging Filters pane, click the filters icon on the right-hand side.

- To clear the current messages, right-click an event and click **Clear Content**.
- To save the current messages to a file on your PC, right-click an event and click **Save Content**.
- To copy the current content, right-click an event and click **Copy**.
- To change the background and foreground colors of syslog messages according to their severity, right-click an event and click **Color Settings**.

Firewall Dashboard Tab

The Firewall Dashboard tab lets you view important information about the traffic passing through your ASA. This dashboard differs depending on whether you are in single context mode or multiple context mode. In multiple context mode, the Firewall Dashboard is viewable within each context.

Figure 4-4 shows some of the elements of the Firewall Dashboard tab.

Figure 4-4 Firewall Dashboard Tab

Legend

GUI Element	Description
1	Traffic Overview Pane, page 4-22
2	Top 10 Access Rules Pane, page 4-23
3	Top Usage Status Pane, page 4-23
(not shown)	Top Ten Protected Servers Under SYN Attack Pane, page 4-24
(not shown)	Top 200 Hosts Pane, page 4-24
(not shown)	Top Botnet Traffic Filter Hits Pane, page 4-24

Traffic Overview Pane

Enabled by default. If you disable basic threat detection (see the "Configuring Basic Threat Detection Statistics" section on page 77-4), then this area includes an Enable button that lets you enable basic threat detection. The runtime statistics include the following information, which is *display-only*:

- The number of connections and NAT translations.
- The rate of dropped packets per second caused by access list denials and application inspections.
- The rate of dropped packets per second that are identified as part of a scanning attack, or that are incomplete sessions detected, such as TCP SYN attack detected or no data UDP session attack detected.

Top 10 Access Rules Pane

Enabled by default. If you disable threat detection statistics for access rules (see the "Configuring Advanced Threat Detection Statistics" section on page 77-5), then this area includes an Enable button that lets you enable statistics for access rules.

In the Table view, you can select a rule in the list and right-click the rule to display a popup menu item, **Show Rule**. Choose this item to go to the Access Rules table and select that rule in this table.

Top Usage Status Pane

Disabled by default. This pane contains the following four tabs:

- Top 10 Services—Threat Detection feature
- Top 10 Sources—Threat Detection feature
- Top 10 Destinations—Threat Detection feature
- Top 10 Users—Identity Firewall feature

The first three tabs—Top 10 Services, Top 10 Sources, and Top 10 Destinations—provide statistics for threat detection features. Each tab includes an Enable button that let you enable each threat detection feature. You can enable them according to the "Configuring Basic Threat Detection Statistics" section on page 77-4.

The Top 10 Services Enable button enables statistics for both ports and protocols (both must be enabled for the display). The Top 10 Sources and Top 10 Destinations Enable buttons enable statistics for hosts. The top usage status statistics for hosts (sources and destinations), and ports and protocols are displayed.

The fourth tab for Top 10 Users provides statistics for the Identity Firewall feature. The Identity Firewall feature provides access control based on users' identities. You can configure access rules and security policies based on user names and user groups name rather than through source IP addresses. The ASA provides this feature by accessing an IP-user mapping database.

The Top 10 Users tab displays data only when you have configured the Identity Firewall feature in the ASA, which includes configuring these additional components—Microsoft Active Directory and Cisco Active Directory (AD) Agent. See the Configuring the Identity Firewall, page 47-9 for information.

Depending on which option you choose, the Top 10 Users tab shows statistics for received EPS packets, sent EPS packets, and sent attacks for the top 10 users. For each user (displayed as *domain\user_name*), the tab displays the average EPS packet, the current EPS packet, the trigger, and total events for that user.

Caution Enabling statistics can affect the ASA 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 a modest effect.

Top Ten Protected Servers Under SYN Attack Pane

Disabled by default. This area includes an Enable button that lets you enable the feature, or you can enable it according to the "Configuring Basic Threat Detection Statistics" section on page 77-4. Statistics for the top ten protected servers under attack are displayed.

For the average rate of attack, the ASA samples the data every 30 seconds over the rate interval (by default 30 minutes).

If there is more than one attacker, then "<various>" displays, followed by the last attacker IP address.

Click **Detail** to view statistics for all servers (up to 1000) instead of just 10 servers. You can also view history sampling data. The ASA samples the number of attacks 60 times during the rate interval, so for the default 30-minute period, statistics are collected every 60 seconds.

Top 200 Hosts Pane

Disabled by default. Shows the top 200 hosts connected through the ASA. Each entry of a host contains the IP address of the host and the number of connections initiated by the host, and is updated every 120 seconds. To enable this display, enter the **hpm topnenable** command.

Top Botnet Traffic Filter Hits Pane

Disabled by default. This area includes links to configure the Botnet Traffic Filter. Reports of the top ten botnet sites, ports, and infected hosts provide a snapshot of the data, and may not match the top ten items since 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.

For more information, see Configuring the Botnet Traffic Filter.

Cluster Dashboard Tab

The Cluster Dashboard tab shows a summary of cluster membership and resource utilization.

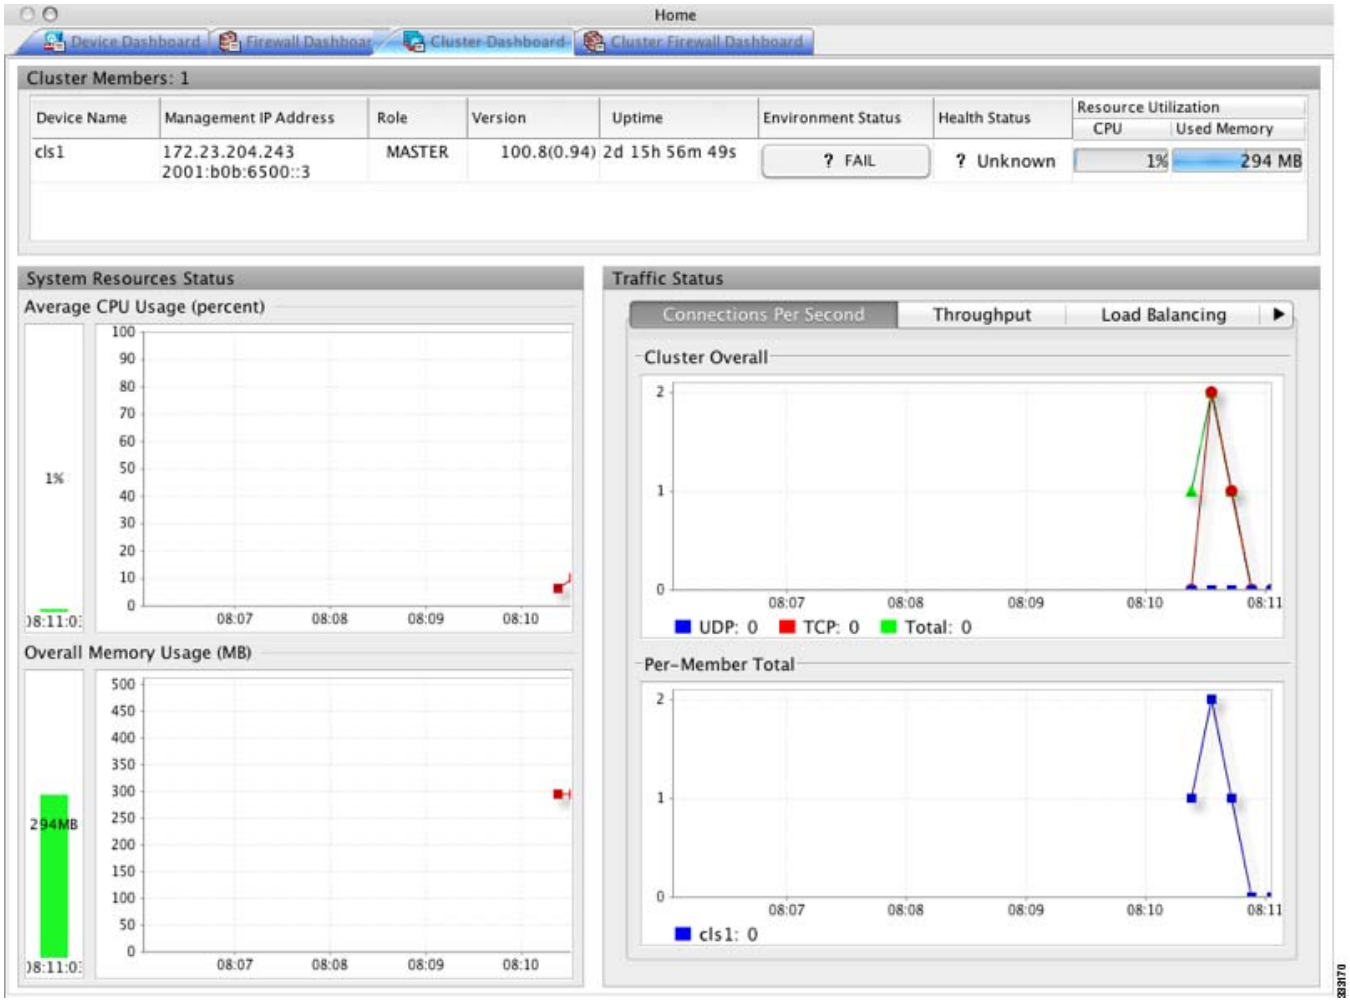

- Cluster Members—Shows the names and basic information about the members comprising the cluster (their management IP address, version, role in the cluster, etc.) and their health status (environment status, Health Status, and resource utilization meters).

  Note In multiple context mode, if you connect ASDM to the admin context, and then change to a different context, the Management IP Address listed does not change to show the current context management IP addresses; it continues to show the admin context management IP addresses, including the main cluster IP address to which ASDM is currently connected.

- System Resource Status—Shows resource utilization (CPU and memory) across the cluster and traffic graphs, both cluster-wide and per-device.
- Traffic Status—Each tab has the following graphs.
  - Connections Per Second tab:

    Cluster Overall—Shows the connections per second throughout the cluster.

Per-Member Total—Shows the average connections per second for each member.

- Throughput tab:

  Cluster Overall—Shows the aggregated egress throughput throughout the cluster.

  Per-Member Throughput—Shows the member throughput, one line per member.

- Load Balacing tab:

  Per-Member Percentage of Total Traffic—For each member, shows the percentage of total cluster traffic that the member receives.

  Per-Member Locally Processed Traffic—For each member, shows the percentage of traffic that was processed locally.

- Control Link Usage tab:

  Per-Member Receival Capacity Utilization—For each member, shows the usage of the transmittal capacity.

  Per-Member Transmittal Capacity Utilization—For each member, shows the usage of the receival capacity.

Cluster Firewall Dashboard Tab

The Cluster Firewall Dashboard tab shows Traffic Overview and the “top N” statistics, similar to those shown in the Firewall Dashboard, but aggregated across the whole cluster.

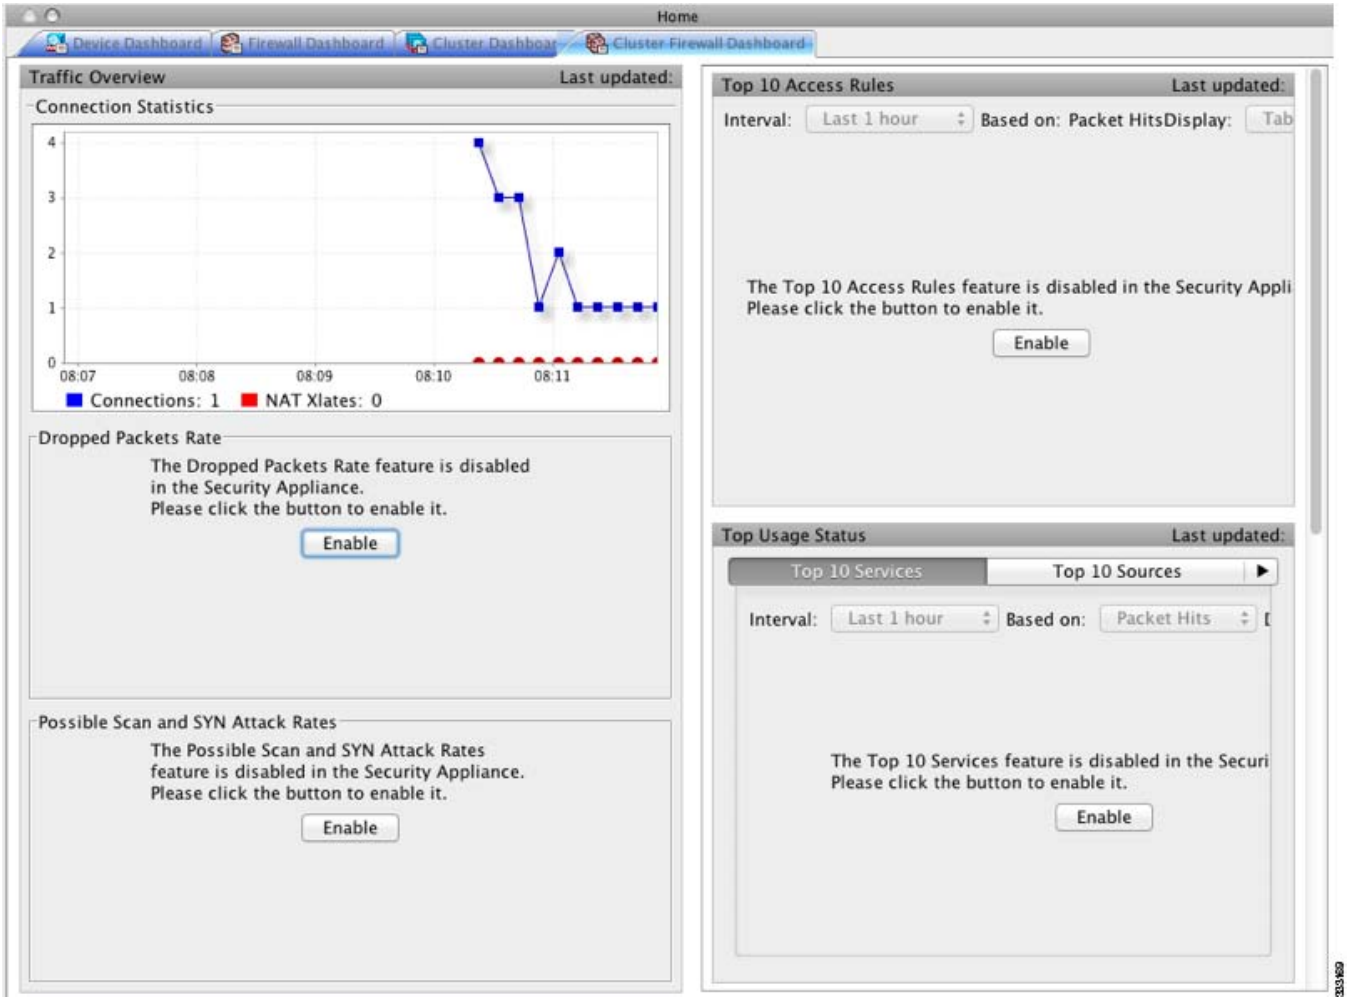

Content Security Tab

The Content Security tab lets you view important information about the Content Security and Control (CSC) SSM. This pane appears only if CSC software running on the CSC SSM is installed in the ASA.

For an introduction to the CSC SSM, see the "Information About the CSC SSM" section on page 83-1.

Note If you have not completed the CSC Setup Wizard by choosing **Configuration > Trend Micro Content Security > CSC Setup**, you cannot access the panes under Home > Content Security. Instead, a dialog box appears and lets you access the CSC Setup Wizard directly from this location.

Figure 4-5 shows the elements of the Content Security tab.

Figure 4-5 Content Security Tab

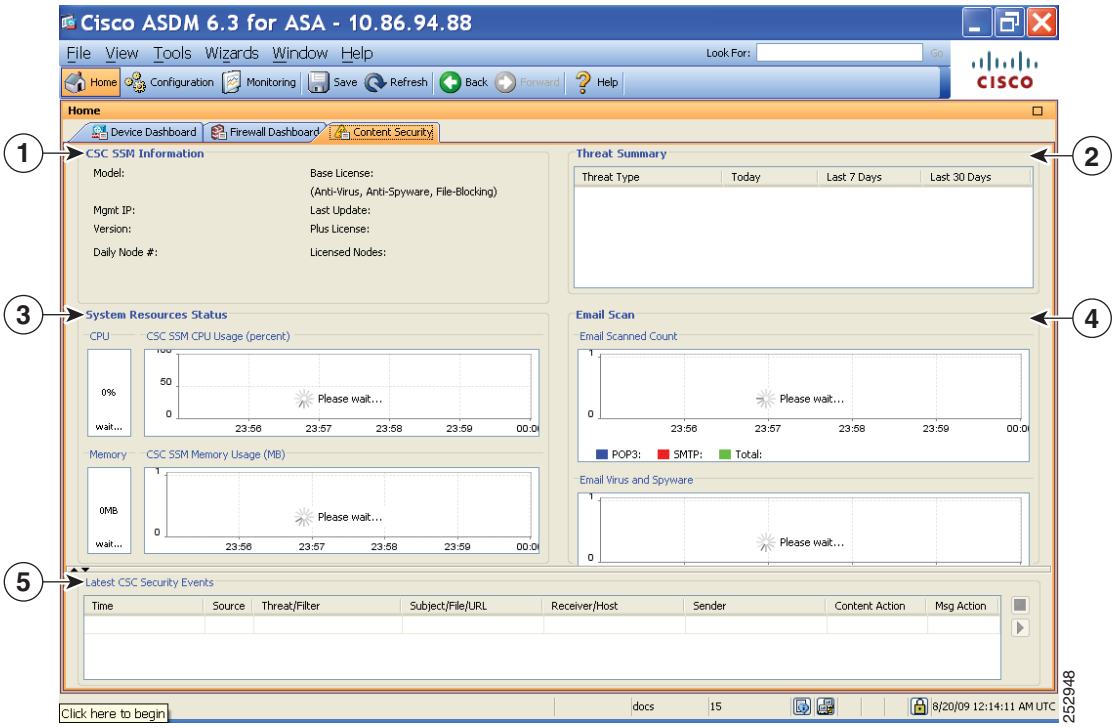

Legend

GUI Element	Description
1	CSC SSM Information pane.
2	Threat Summary pane. Shows aggregated data about threats detected by the CSC SSM, including the following threat types: Virus, Spyware, URL Filtered or Blocked, Spam. Blocked, Files Blocked, and Damage Control Services.
3	System Resources Status pane.
4	Email Scan pane. The graphs display data in ten-second intervals.
5	Latest CSC Security Events pane.

Intrusion Prevention Tab

The Intrusion Prevention tab lets you view important information about IPS. This tab appears only when you have an IPS module installed on the ASA.

To connect to the IPS module, perform the following steps:

Step 1 In the main ASDM application window, click the **Intrusion Prevention** tab.

The Connecting to IPS dialog box appears.

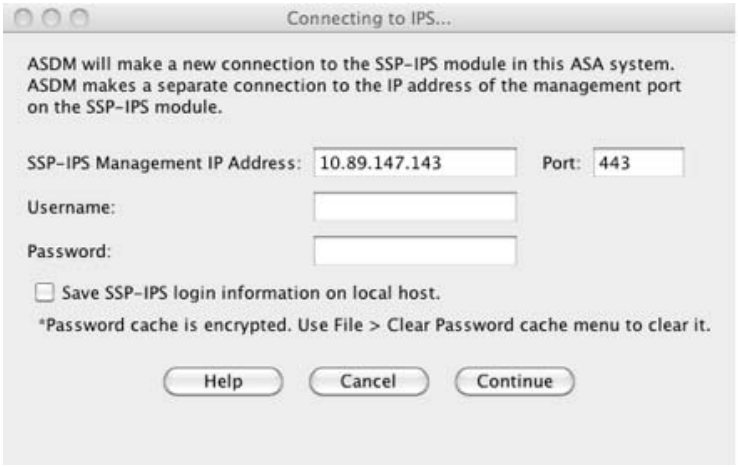

Step 2 Enter the IP address, port, username and password.The default IP address and port is 192.168.1.2:443. The default username and password is **cisco** and **cisco**

Step 3 To save the login information on your local PC, check the **Save IPS login** information on local host check box.

Step 4 Click **Continue**.

For more information about intrusion prevention, see Chapter 81, "Configuring the ASA IPS Module.".

Figure 4-6 shows the elements of the Health Dashboard tab, located on the Intrusion Prevention tab.

Figure 4-6 Intrusion Prevention Tab (Health Dashboard)

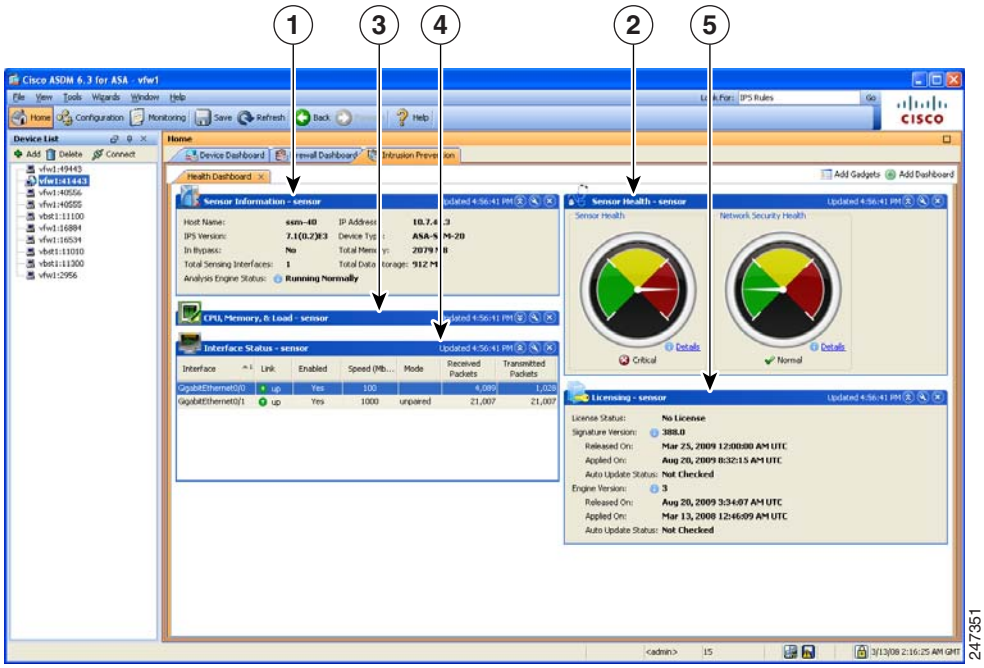

Legend

GUI Element	Description
1	Sensor Information pane.
2	Sensor Health pane.
3	CPU, Memory, and Load pane.
4	Interface Status pane.
5	Licensing pane.

ASA CX Status Tab

The ASA CX Status tab lets you view important information about the ASA CX module. This tab appears only when you have an ASA CX module installed on the ASA.

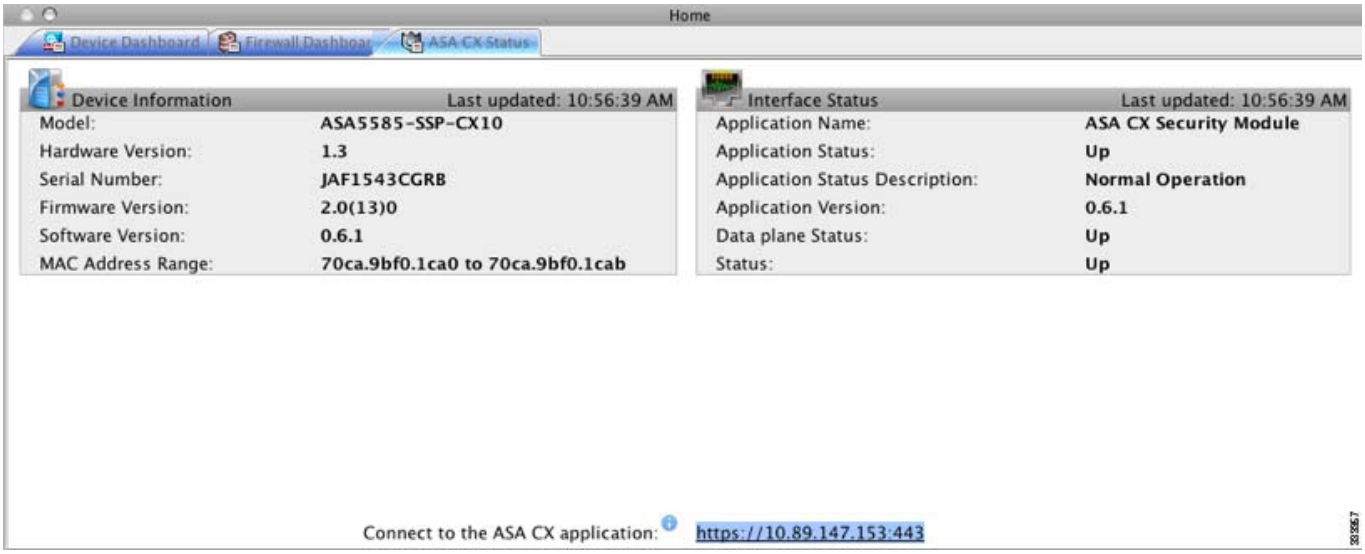

Home Pane (System)

The ASDM System Home pane lets you view important status information about your ASA. Many of the details available in the ASDM System Home pane are available elsewhere in ASDM, but this pane shows at-a-glance how your ASA is running. Status information in the System Home pane is updated every ten seconds.

Figure 4-7 on page 4-31 shows the elements of the System Home pane.

Figure 4-7 System Home Pane

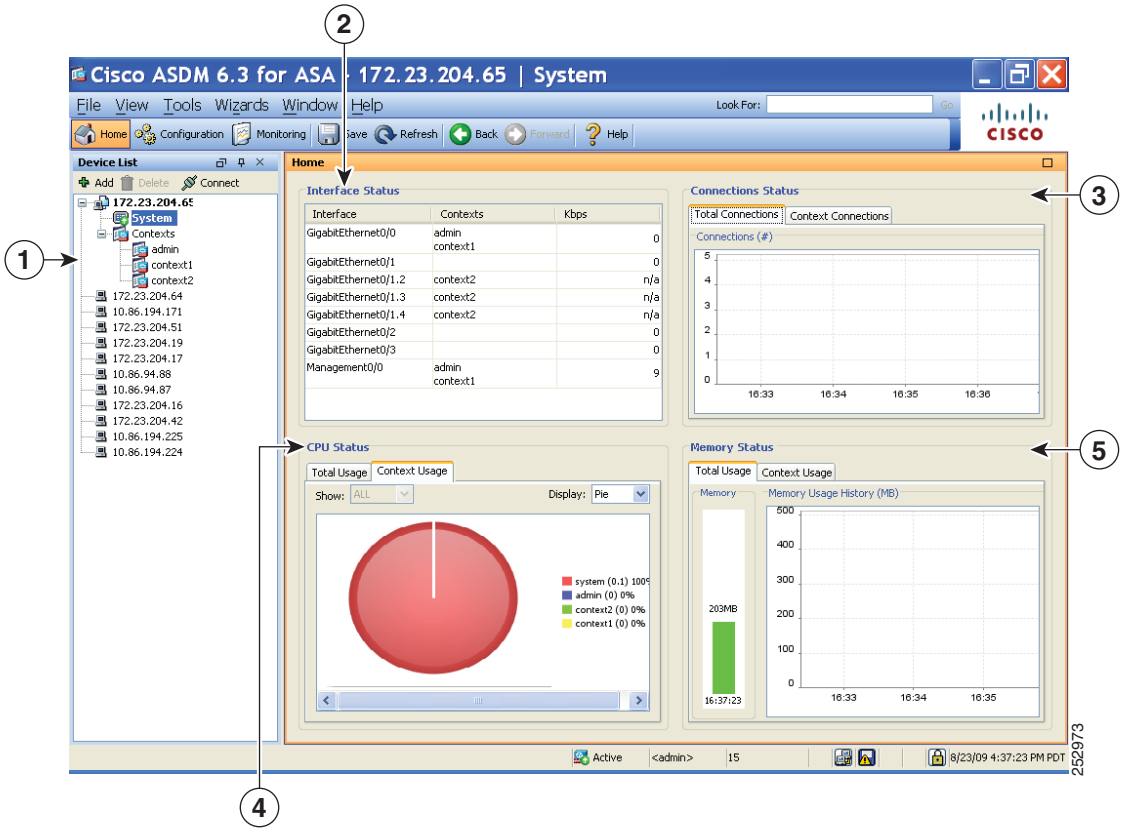

Legend

GUI Element	Description
1	System vs. Context selection.
2	Interface Status pane. Choose an interface to view the total amount of traffic through the interface.
3	Connection Status pane.
4	CPU Status pane.
5	Memory Status pane.

Defining ASDM Preferences

This feature lets you define the behavior of certain ASDM settings.

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 Table** 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 **Warn of insufficient ASA memory when ASDM loads** check box to receive notification when the minimum amount of ASA memory is insufficient to run complete functionality in the ASDM application. ASDM displays the memory warning in a text banner message at bootup, displays a message in the title bar text in ASDM, and sends a syslog alert once every 24 hours.

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

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

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

i. 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 Table 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 ASA 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 ASA 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.

Using the ASDM Assistant

The ASDM Assistant tool lets you search and view useful ASDM procedural help about certain tasks.

To access information, choose **View > ASDM Assistant > How Do I?** or enter a search request from the Look For field in the menu bar. From the Find drop-down list, choose **How Do I?** to begin the search.

Note This feature is not available on the PIX security appliance.

To view the ASDM Assistant, perform the following steps:

Step 1 In the main ASDM application window, choose **View > ASDM Assistant**.

The ASDM Assistant pane appears.

Step 2 In the Search field, enter the information that you want to find, and click **Go**.

The requested information appears in the Search Results pane.

Step 3 Click any links that appear in the Search Results and Features sections to obtain more details.

Enabling History Metrics

The Configuration > Device Management > Advanced > History Metrics pane lets you configure the adaptive ASA 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**.

Unsupported Commands

ASDM supports almost all commands available for the adaptive ASA, 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 4-35
- Effects of Unsupported Commands, page 4-35
- Discontinuous Subnet Masks Not Supported, page 4-36
- Interactive User Commands Not Supported by the ASDM CLI Tool, page 4-36

Ignored and View-Only Commands

Table 4-5 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 4-5 List of Unsupported Commands

Unsupported Commands	ASDM Behavior
capture	Ignored.
coredump	Ignored. This can be configured only using the CLI.
crypto engine large-mod-accel	Ignored.
dhcp-server (tunnel-group name general-attributes)	ASDM only allows one setting for all DHCP servers.
eject	Unsupported.
established	Ignored.
failover timeout	Ignored.
fips	Ignored.
nat-assigned-to-public-ip	Ignored.
pager	Ignored.
pim accept-register route-map	Ignored. You can configure only the **list** option using ASDM.

service-policy global	Ignored if it uses a **match access-list** class. For example:
	`access-list myacl extended permit ip any any`
	`class-map mycm`
	`  match access-list myacl`
	`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.
threat-detection rate	Ignored.

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**.

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. Choose **Tools > 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

Managing Feature Licenses

A license specifies the options that are enabled on a given ASA. This document describes how to obtain a license activation key and how to activate it. It also describes the available licenses for each model.

Note This chapter describes licensing for Version 7.0; for other versions, see the licensing documentation that applies to your version:

http://www.cisco.com/en/US/products/ps6120/products_licensing_information_listing.html

This chapter includes the following sections:

- Supported Feature Licenses Per Model, page 5-1
- Information About Feature Licenses, page 5-23
- Guidelines and Limitations, page 5-34
- Configuring Licenses, page 5-36
- Monitoring Licenses, page 5-40
- Feature History for Licensing, page 5-41

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 5-1
- License Notes, page 5-18
- VPN License and Feature Compatibility, page 5-23

Licenses Per Model

This section lists the feature licenses available for each model:

- ASA 5505, page 5-2
- ASA 5510, page 5-3
- ASA 5520, page 5-4

- ASA 5540, page 5-5
- ASA 5550, page 5-6
- ASA 5580, page 5-7
- ASA 5512-X, page 5-8
- ASA 5515-X, page 5-9
- ASA 5525-X, page 5-10
- ASA 5545-X, page 5-11
- ASA 5555-X, page 5-12
- ASA 5585-X with SSP-10, page 5-13
- ASA 5585-X with SSP-20, page 5-14
- ASA 5585-X with SSP-40 and -60, page 5-15
- ASA Services Module, page 5-16

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 24 Unified Communications license plus the Strong Encryption license; or the 500 AnyConnect Premium license plus the GTP/GPRS license; or all four licenses together.

For detailed information about licenses, see the "License Notes" section on page 5-18.

ASA 5505

Table 5-1 ASA 5505 License Features

Licenses	Description (Base License in Plain Text)	Description (Security Plus Lic. in Plain Text)
Firewall Licenses		
Botnet Traffic Filter	Disabled	*Opt. Time-based lic: Available*	Disabled	*Opt. Time-based lic: Available*
Firewall Conns, Concurrent	10,000	25,000
GTP/GPRS	No support	No support
Intercompany Media Eng.	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
UC Phone Proxy Sessions	2	*Optional license: 24*	2	*Optional license: 24*
VPN Licenses		
Adv. Endpoint Assessment	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect Essentials	Disabled	*Optional license: Available (25 sessions)*	Disabled	*Optional license: Available (25 sessions)*
AnyConnect for Mobile	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:* *10* *25*	2	*Optional Permanent or Time-based licenses:* *10* *25*
Other VPN (sessions)	10	25
Total VPN (sessions), combined all types	up to 25[1]	up to 25
VPN Load Balancing	No support	No support

Table 5-1 ASA 5505 License Features (continued)

Licenses	Description (Base License in Plain Text)				Description (Security Plus Lic. in Plain Text)			
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)			
Interfaces of all types, Max.	52				120			
Security Contexts	No support				No support			
Clustering	No support				No support			
Inside Hosts, concurrent[2]	10[3]	*Opt. licenses:*	*50*	*Unlimited*	10[3]	*Opt. licenses:*	*50*	*Unlimited*
VLANs, maximum	Routed mode: 3 (2 regular and 1 restricted) Transparent mode: 2				Routed mode: 20 Transparent mode: 3 (2 regular and 1 failover)			
VLAN Trunks, maximum	No support				8 trunks			

1. The total number of VPN sessions depends on your licenses. If you enable AnyConnect Essentials, then the total is the model maximum of 25. If you enable AnyConnect Premium, then the total is the AnyConnect Premium value plus the Other VPN value, not to exceed 25 sessions.
2. In routed mode, hosts on the inside (Business and Home VLANs) count toward the limit when they communicate with the outside (Internet VLAN), including when the inside initiates a connection to the outside as well as when the outside initiates a connection to the inside. Note that even when the outside initiates a connection to the inside, outside hosts are *not* counted toward the limit; only the inside hosts count. Hosts that initiate traffic between Business and Home are also not counted toward the limit. The interface associated with the default route is considered to be the outside 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 toward the host limit. Use 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.

ASA 5510

Table 5-2 ASA 5510 License Features

Licenses	Description (Base License in Plain Text)				Description (Security Plus Lic. in Plain Text)			
Firewall Licenses								
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*			Disabled	*Optional Time-based license: Available*		
Firewall Conns, Concurrent	50,000				130,000			
GTP/GPRS	No support				No support			
Intercompany Media Eng.	Disabled	*Optional license: Available*			Disabled	*Optional license: Available*		
UC Phone Proxy Sessions	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 for Cisco VPN Phone	Disabled	*Optional license: Available*			Disabled	*Optional license: Available*		
AnyConnect Essentials	Disabled	*Optional license: Available (250 sessions)*			Disabled	*Optional license: Available (250 sessions)*		
AnyConnect for Mobile	Disabled	*Optional license: Available*			Disabled	*Optional license: Available*		

Table 5-2 ASA 5510 License Features (continued)

Licenses	Description (Base License in Plain Text)						Description (Security Plus Lic. in Plain Text)					
AnyConnect Premium (sessions)	2	*Optional Perm. or Time-based lic,:*					2	*Optional Perm. or Time-based lic:*				
		10	*25*	*50*	*100*	*250*		*10*	*25*	*50*	*100*	*250*
	Optional Shared licenses: Participant or Server. For the Server:						*Optional Shared licenses: Participant or Server. For the Server:*					
	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*		
Total VPN (sessions), combined all types	250						250					
Other VPN (sessions)	250						250					
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					
Interfaces of all types, Max.	240						440					
Interface Speed	All: Fast Ethernet						Ethernet 0/0 and 0/1: Gigabit Ethernet[1] Ethernet 0/2, 0/3, 0/4 (and others): Fast Eth.					
Security Contexts	No support						2	*Optional licenses:*				*5*
Clustering	No support						No support					
VLANs, Maximum	50						100					

1. Although the Ethernet 0/0 and 0/1 ports are Gigabit Ethernet, they are still identified as "Ethernet" in the software.

ASA 5520

Table 5-3 ASA 5520 License Features

Licenses	Description (Base License in Plain Text)								
Firewall Licenses									
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*							
Firewall Conns, Concurrent	280,000								
GTP/GPRS	Disabled	*Optional license: Available*							
Intercompany Media Eng.	Disabled	*Optional license: Available*							
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	*100*	*250*	*500*	*750*	*1000*
VPN Licenses									
Adv. Endpoint Assessment	Disabled	*Optional license: Available*							
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*							
AnyConnect Essentials	Disabled	*Optional license: Available (750 sessions)*							
AnyConnect for Mobile	Disabled	*Optional license: Available*							

Table 5-3 ASA 5520 License Features (continued)

Licenses	Description (Base License in Plain Text)							
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*						
		10	*25*	*50*	*100*	*250*	*500*	*750*
	Optional Shared licenses: Participant or Server. For the Server:							
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*		
Total VPN (sessions), combined all types	750							
Other VPN (sessions)	750							
VPN Load Balancing	Supported							
General Licenses								
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*						
Failover	Active/Standby or Active/Active							
Interfaces of all types, Max.	640							
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*			
Clustering	No support							
VLANs, Maximum	150							

ASA 5540

Table 5-4 ASA 5540 License Features

Licenses	Description (Base License in Plain Text)									
Firewall Licenses										
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*								
Firewall Conns, Concurrent	400,000									
GTP/GPRS	Disabled	*Optional license: Available*								
Intercompany Media Eng.	Disabled	*Optional license: Available*								
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*
VPN Licenses										
Adv. Endpoint Assessment	Disabled	*Optional license: Available*								
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*								
AnyConnect Essentials	Disabled	*Optional license: Available (2500 sessions)*								
AnyConnect for Mobile	Disabled	*Optional license: Available*								
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*								
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*
	Optional Shared licenses: Participant or Server. For the Server:									
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*				
Total VPN (sessions), combined all types	5000									

Table 5-4 ASA 5540 License Features (continued)

Licenses	Description (Base License in Plain Text)					
Other VPN (sessions)	5000					
VPN Load Balancing	Supported					
General Licenses						
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*				
Failover	Active/Standby or Active/Active					
Interfaces of all types, Max.	840					
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*	*50*
Clustering	No support					
VLANs, Maximum	200					

ASA 5550

Table 5-5 ASA 5550 License Features

Licenses	Description (Base License in Plain Text)								
Firewall Licenses									
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*							
Firewall Conns, Concurrent	650,000								
GTP/GPRS	Disabled	*Optional license: Available*							
Intercompany Media Eng.	Disabled	*Optional license: Available*							
UC Phone Proxy Sessions	2	*Optional licenses:*							
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*
VPN Licenses									
Adv. Endpoint Assessment	Disabled	*Optional license: Available*							
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*							
AnyConnect Essentials	Disabled	*Optional license: Available (5000 sessions)*							
AnyConnect for Mobile	Disabled	*Optional license: Available*							
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*							
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*
	Optional Shared licenses: Participant or Server. For the Server:								
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*			
Total VPN (sessions), combined all types	5000								
Other VPN (sessions)	5000								
VPN Load Balancing	Supported								
General Licenses									
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*							

Table 5-5 ASA 5550 License Features (continued)

Licenses	Description (Base License in Plain Text)						
Failover	Active/Standby or Active/Active						
Interfaces of all types, Max.	1640						
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*	*50*	*100*
Clustering	No support						
VLANs, Maximum	400						

ASA 5580

Table 5-6 ASA 5580 License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*									
Firewall Conns, Concurrent	5580-20: 2,000,000					5580-40: 4,000,000					
GTP/GPRS	Disabled	*Optional license: Available*									
Intercompany Media Eng.	Disabled	*Optional license: Available*									
UC Phone Proxy Sessions	2	*Optional licenses:*									
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*	*5000*	*10,000*[1]
VPN Licenses											
Adv. Endpoint Assessment	Disabled	*Optional license: Available*									
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*									
AnyConnect Essentials	Disabled	*Optional license: Available (10000 sessions)*									
AnyConnect for Mobile	Disabled	*Optional license: Available*									
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*									
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*	*10,000*
	Optional Shared licenses: Participant or Server. For the Server:										
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*					
Total VPN (sessions), combined all types	10,000										
Other VPN (sessions)	10,000										
VPN Load Balancing	Supported										
General Licenses											
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*									
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	4176										
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*	*250*		

Table 5-6 ASA 5580 License Features (continued)

Licenses	Description (Base License in Plain Text)	
Clustering	Disabled	*Optional license: Available*
VLANs, Maximum	1024	

1. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

ASA 5512-X

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-7 ASA 5512-X License Features

Licenses	Description (Base License in Plain Text)					Description (Security Plus Lic. in Plain Text)				
Firewall Licenses										
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*				Disabled	*Optional Time-based license: Available*			
Firewall Conns, Concurrent	100,000					250,000				
GTP/GPRS	No support					Disabled	*Optional license: Available*			
Intercompany Media Eng.	Disabled	*Optional license: Available*				Disabled	*Optional license: Available*			
UC Phone Proxy Sessions	2	*Optional licenses:*				2	*Optional licenses:*			
		24	*50*	*100*	*250*		*24*	*50*	*100*	*250*
VPN Licenses										
Adv. Endpoint Assessment	Disabled	*Optional license: Available*				Disabled	*Optional license: Available*			
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*				Disabled	*Optional license: Available*			
AnyConnect Essentials	Disabled	*Optional license: Available (250 sessions)*				Disabled	*Optional license: Available (250 sessions)*			
AnyConnect for Mobile	Disabled	*Optional license: Available*				Disabled	*Optional license: Available*			
AnyConnect Premium (sessions)	2	*Optional Perm. or Time-based lic,:*				2	*Optional Perm. or Time-based lic:*			
		10	*25*	*50*	*100*	*250*	*10*	*25*	*50*	*100*	*250*
	Optional Shared licenses: Participant or Server. For the Server:					*Optional Shared licenses: Participant or Server. For the Server:*				
	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*		
Total VPN (sessions), combined all types	250					250				
Other VPN (sessions)	250					250				
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				
Interfaces of all types, Max.	328					528				

Table 5-7 ASA 5512-X License Features (continued)

Licenses	Description (Base License in Plain Text)		Description (Security Plus Lic. in Plain Text)		
Security Contexts	No support		2	*Optional licenses:*	*5*
IPS Module	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*	
VLANs, Maximum	50		100		

ASA 5515-X

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-8 ASA 5515-X License Features

Licenses	Description (Base License in Plain Text)						
Firewall Licenses							
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*					
Firewall Conns, Concurrent	250,000						
GTP/GPRS	Disabled	*Optional license: Available*					
Intercompany Media Eng.	Disabled	*Optional license: Available*					
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	*100*	*250*	*500*
VPN Licenses							
Adv. Endpoint Assessment	Disabled	*Optional license: Available*					
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*					
AnyConnect Essentials	Disabled	*Optional license: Available (250 sessions)*					
AnyConnect for Mobile	Disabled	*Optional license: Available*					
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*					
		10	*25*	*50*	*100*	*250*	
	Optional Shared licenses: Participant or Server. For the Server:						
	500-50,000 in increments of 500				*50,000-545,000 in increments of 1000*		
Total VPN (sessions), combined all types	250						
Other VPN (sessions)	250						
VPN Load Balancing	Supported						
General Licenses							
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*					
Failover	Active/Standby or Active/Active						
Interfaces of all types, Max.	528						
Security Contexts	2	*Optional licenses:*	*5*				
Clustering	No support						
IPS Module	Disabled	*Optional license: Available*					
VLANs, Maximum	100						

ASA 5525-X

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-9 ASA 5525-X License Features

Licenses	Description (Base License in Plain Text)								
Firewall Licenses									
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*							
Firewall Conns, Concurrent	500,000								
GTP/GPRS	Disabled	*Optional license: Available*							
Intercompany Media Eng.	Disabled	*Optional license: Available*							
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	*100*	*250*	*500*	*750*	*1000*
VPN Licenses									
Adv. Endpoint Assessment	Disabled	*Optional license: Available*							
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*							
AnyConnect Essentials	Disabled	*Optional license: Available (750 sessions)*							
AnyConnect for Mobile	Disabled	*Optional license: Available*							
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*							
		10	*25*	*50*	*100*	*250*	*500*	*750*	
	Optional Shared licenses: Participant or Server. For the Server:								
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*			
Total VPN (sessions), combined all types	750								
Other VPN (sessions)	750								
VPN Load Balancing	Supported								
General Licenses									
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*							
Failover	Active/Standby or Active/Active								
Interfaces of all types, Max.	928								
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*				
Clustering	No support								
IPS Module	Disabled	*Optional license: Available*							
VLANs, Maximum	200								

ASA 5545-X

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-10 ASA 5545-X License Features

Licenses	Description (Base License in Plain Text)									
Firewall Licenses										
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*								
Firewall Conns, Concurrent	750,000									
GTP/GPRS	Disabled	*Optional license: Available*								
Intercompany Media Eng.	Disabled	*Optional license: Available*								
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*
VPN Licenses										
Adv. Endpoint Assessment	Disabled	*Optional license: Available*								
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*								
AnyConnect Essentials	Disabled	*Optional license: Available (2500 sessions)*								
AnyConnect for Mobile	Disabled	*Optional license: Available*								
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*								
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*
	Optional Shared licenses: Participant or Server. For the Server:									
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*				
Total VPN (sessions), combined all types	2500									
Other VPN (sessions)	2500									
VPN Load Balancing	Supported									
General Licenses										
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*								
Failover	Active/Standby or Active/Active									
Interfaces of all types, Max.	1328									
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*	*50*				
Clustering	No support									
IPS Module	Disabled	*Optional license: Available*								
VLANs, Maximum	300									

ASA 5555-X

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-11 ASA 5555-X License Features

Licenses	Description (Base License in Plain Text)									
Firewall Licenses										
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*								
Firewall Conns, Concurrent	1,000,000									
GTP/GPRS	Disabled	*Optional license: Available*								
Intercompany Media Eng.	Disabled	*Optional license: Available*								
UC Phone Proxy Sessions	2	*Optional licenses:*								
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*
VPN Licenses										
Adv. Endpoint Assessment	Disabled	*Optional license: Available*								
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*								
AnyConnect Essentials	Disabled	*Optional license: Available (5000 sessions)*								
AnyConnect for Mobile	Disabled	*Optional license: Available*								
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*								
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*
	Optional Shared licenses: Participant or Server. For the Server:									
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*				
Total VPN (sessions), combined all types	5000									
Other VPN (sessions)	5000									
VPN Load Balancing	Supported									
General Licenses										
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*								
Failover	Active/Standby or Active/Active									
Interfaces of all types, Max.	2128									
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*		
Clustering	No support									
IPS Module	Disabled	*Optional license: Available*								
VLANs, Maximum	500									

ASA 5585-X with SSP-10

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-12 ASA 5585-X with SSP-10 License Features

Licenses	Description (Base License in Plain Text)	Description (Security Plus License in Plain Text)
Firewall Licenses		
Botnet Traffic Filter	Disabled	*Opt. Time-based lic: Available*	Disabled	*Opt. Time-based lic: Available*
Firewall Conns, Concurrent	1,000,000	1,000,000
GTP/GPRS	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
Intercompany Media Eng.	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
UC Phone Proxy Sessions	2	*Optional licenses:*	2	*Optional licenses:*
	24	*50*	*100*	*250*	*500*	*24*	*50*	*100*	*250*	*500*
	750	*1000*	*2000*	*3000*	*750*	*1000*	*2000*	*3000*
VPN Licenses[1]		
Adv. Endpoint Assessment	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect Essentials	Disabled	*Optional license: Available (5000 sessions)*	Disabled	*Optional license: Available (5000 sessions)*
AnyConnect for Mobile	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect Premium (sessions)	2	*Opt. Permanent or Time-based lic.:*	2	*Opt. Permanent or Time-based lic.:*
	10	*25*	*50*	*100*	*250*	*10*	*25*	*50*	*100*	*250*
	500	*750*	*1000*	*2500*	*5000*	*500*	*750*	*1000*	*2500*	*5000*
	Optional Shared licenses: Participant or Server. For the Server:	*Optional Shared licenses: Participant or Server. For the Server:*
	500-50,000 in increments of 500	*500-50,000 in increments of 500*
	50,000-545,000 in increments of 1000	*50,000-545,000 in increments of 1000*
Total VPN (sessions), combined all types	5000	5000
Other VPN (sessions)	5000	5000
VPN Load Balancing	Supported	Supported
General Licenses		
10 GE I/O	Disabled; fiber ifcs run at 1 GE	Enabled; fiber ifcs run at 10 GE
Encryption	Base (DES)	*Opt. lic.: Strong (3DES/AES)*	Base (DES)	*Opt. lic.: Strong (3DES/AES)*
Failover	Active/Standby or Active/Active	Active/Standby or Active/Active
Interfaces of all types, Max.	4176	4176
Security Contexts	2	*Optional licenses:*	2	*Optional licenses:*
	5	*10*	*20*	*50*	*100*	*5*	*10*	*20*	*50*	*100*

Table 5-12 ASA 5585-X with SSP-10 License Features (continued)

Licenses	Description (Base License in Plain Text)	Description (Security Plus License in Plain Text)
Clustering	Disabled	*Optional license: Available*
VLANs, Maximum	1024	1024

1. See the "VPN License and Feature Compatibility" section on page 5-23.

ASA 5585-X with SSP-20

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-13 ASA 5585-X with SSP-20 License Features

Licenses	Description (Base License in Plain Text)	Description (Security Plus Lic. in Plain Text)
Firewall Licenses		
Botnet Traffic Filter	Disabled	*Opt. Time-based lic: Available*	Disabled	*Opt. Time-based lic: Available*
Firewall Conns, Concurrent	2,000,000	2,000,000
GTP/GPRS	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
Intercompany Media Eng.	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
UC Phone Proxy Sessions	2	*Optional licenses:*	*24*	*50*	2	*Optional licenses:*	*24*	*50*
		100	*250*	*500*	*750*	*1000*		*100*	*250*	*500*	*750*	*1000*
		2000	*3000*	*5000*	10,000[1]		*2000*	*3000*	*5000*	10,000[1]
VPN Licenses		
Adv. Endpoint Assessment	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*
AnyConnect Essentials	Disabled	*Optional license: Available (10,000 sessions)*	Disabled	*Optional license: Available (10,000 sessions)*
AnyConnect for Mobile	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*	*10*	*25*	2	*Optional Permanent or Time-based licenses:*	*10*	*25*
		50	*100*	*250*	*500*	*750*		*50*	*100*	*250*	*500*	*750*
		1000	*2500*	*5000*	10,000		*1000*	*2500*	*5000*	10,000	
	Optional Shared licenses: Participant or Server. For the Server:	*Optional Shared licenses: Participant or Server. For the Server:*
	500-50,000 in increments of 500	*500-50,000 in increments of 500*
	50,000-545,000 in increments of 1000	*50,000-545,000 in increments of 1000*
Total VPN (sessions), combined all types	10,000	10,000
Other VPN (sessions)	10,000	10,000
VPN Load Balancing	Supported	Supported
General Licenses		

Table 5-13 ASA 5585-X with SSP-20 License Features (continued)

Licenses	Description (Base License in Plain Text)						Description (Security Plus Lic. in Plain Text)					
10 GE I/O	Disabled; fiber ifcs run at 1 GE						Enabled; fiber ifcs run at 10 GE					
Encryption	Base (DES)	*Opt. lic.: Strong (3DES/AES)*					Base (DES)	*Opt. lic.: Strong (3DES/AES)*				
Failover	Active/Standby or Active/Active						Active/Standby or Active/Active					
Interfaces of all types, Max.	4176						4176					
Security Contexts	2	*Optional licenses:*			*5*	*10*	2	*Optional licenses:*			*5*	*10*
		20	*50*	*100*	250			*20*	*50*	*100*	250	
Clustering	Disabled	*Optional license: Available*										
VLANs, Maximum	1024						1024					

1. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

ASA 5585-X with SSP-40 and -60

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Note For SSP-40 and SSP-60, you can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired. When using two SSPs in the chassis, VPN is not supported; note, however, that VPN has not been disabled.

Table 5-14 ASA 5585-X with SSP-40 and -60 License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*									
Firewall Conns, Concurrent	5585-X with SSP-40: 4,000,000					5585-X with SSP-60: 10,000,000					
GTP/GPRS	Disabled	*Optional license: Available*									
Intercompany Media Eng.	Disabled	*Optional license: Available*									
UC Phone Proxy Sessions	2	*Optional licenses:*									
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*	*5000*	*10,000*[1]
VPN Licenses											
Adv. Endpoint Assessment	Disabled	*Optional license: Available*									
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*									
AnyConnect Essentials	Disabled	*Optional license: Available (10,000 sessions)*									
AnyConnect for Mobile	Disabled	*Optional license: Available*									

Table 5-14 ASA 5585-X with SSP-40 and -60 License Features (continued)

Licenses	Description (Base License in Plain Text)										
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*									
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*	*10,000*
	Optional Shared licenses: Participant or Server. For the Server:										
	500-50,000 in increments of 500						*50,000-545,000 in increments of 1000*				
Total VPN (sessions), combined all types	10,000										
Other VPN (sessions)	10,000										
VPN Load Balancing	Supported										
General Licenses											
10 GE I/O	Enabled; fiber ifcs run at 10 GE										
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*									
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	4176										
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*	*50*	*100*	*250*			
Clustering	Disabled	*Optional license: Available*									
VLANs, Maximum	1024										

1. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

ASA Services Module

If you have a No Payload Encryption model, then some of the features below are not supported. See the "No Payload Encryption Models" section on page 5-33 for a list of unsupported features.

Table 5-15 ASASM License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*									
Firewall Conns, Concurrent	10,000,000										
GTP/GPRS	Disabled	*Optional license: Available*									
Intercompany Media Eng.	Disabled	*Optional license: Available*									
UC Phone Proxy Sessions	2	*Optional licenses:*									
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*	*5000*	*10,000*[1]
VPN Licenses											
Adv. Endpoint Assessment	Disabled	*Optional license: Available*									
AnyConnect for Cisco VPN Phone	Disabled	*Optional license: Available*									
AnyConnect Essentials	Disabled	*Optional license: Available (10,000 sessions)*									
AnyConnect for Mobile	Disabled	*Optional license: Available*									

Table 5-15 ASASM License Features (continued)

Licenses	Description (Base License in Plain Text)										
AnyConnect Premium (sessions)	2	*Optional Permanent or Time-based licenses:*									
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*	*10,000*
	Optional Shared licenses: Participant or Server. For the Server:										
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*					
Total VPN (sessions), combined all types	10,000										
Other VPN (sessions)	10,000										
VPN Load Balancing	Supported										
General Licenses											
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*									
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	4224										
Security Contexts	2	*Optional licenses:*									
		5	*10*	*20*	*50*	*100*	*250*				
Clustering	No support										
VLANs, Maximum	1024										

1. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

License Notes

Table 5-16 includes common footnotes shared by multiple tables in the "Licenses Per Model" section on page 5-1.

Table 5-16 License Notes

License	Notes
AnyConnect Essentials	AnyConnect Essentials sessions include the following VPN types: • SSL VPN • IPsec remote access VPN using IKEv2 This license does not support browser-based (clientless) SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium 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 license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given ASA: AnyConnect Premium license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium licenses on different ASAs in the same network. By default, the ASA uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the **no anyconnect-essentials** command or in ASDM, using the **Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials** pane. See also the "VPN License and Feature Compatibility" section on page 5-23.
AnyConnect for Cisco VPN Phone	In conjunction with an AnyConnect Premium license, this license enables access from hardware IP phones that have built in AnyConnect compatibility.

Table 5-16 License Notes (continued)

License	Notes
AnyConnect for 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 of SSL VPN sessions permitted: AnyConnect Essentials or AnyConnect Premium. **Mobile Posture Support** Enforcing remote access controls and gathering posture data from mobile devices requires an AnyConnect Mobile license and either an AnyConnect Essentials or AnyConnect Premium license to be installed on the ASA. Here is the functionality you receive based on the license you install. • AnyConnect Premium License Functionality – Enforce DAP policies on supported mobile devices based on DAP attributes and any other existing endpoint attributes. This includes allowing or denying remote access from a mobile device. • AnyConnect Essentials License Functionality – Enable or disable mobile device access on a per group basis and to configure that feature using ASDM. – Display information about connected mobile devices via CLI or ASDM without having the ability to enforce DAP policies or deny or allow remote access to those mobile devices.
AnyConnect Premium	AnyConnect Premium sessions include the following VPN types: • SSL VPN • Clientless SSL VPN • IPsec remote access VPN using IKEv2
AnyConnect Premium Shared	A shared license lets the ASA act as a shared license server for multiple client ASAs. The shared license pool is large, but the maximum number of sessions used by each individual ASA cannot exceed the maximum number listed for permanent licenses.
Botnet Traffic Filter	Requires a Strong Encryption (3DES/AES) License to download the dynamic database.
Encryption	The DES license cannot be disabled. If you have the 3DES license installed, DES is still available. To prevent the use of DES when you want to only use strong encryption, be sure to configure any relevant commands to use only string encryption.
Failover, Active/Active	You cannot use Active/Active failover and VPN; if you want to use VPN, use Active/Standby failover.

Table 5-16 License Notes (continued)

License	Notes
Intercompany Media Engine	When you enable the Intercompany Media Engine (IME) license, you can use TLS proxy sessions up to the configured TLS proxy limit. If you also have a Unified Communications (UC) license installed that is higher than the default TLS proxy limit, then the ASA sets the limit to be the UC license limit plus an additional number of sessions depending on your model. You can manually configure the TLS proxy limit using the **tls-proxy maximum-sessions** command or in ASDM, using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. To view the limits of your model, enter the **tls-proxy maximum-sessions ?** command. If you also install the UC license, then the TLS proxy sessions available for UC are also available for IME sessions. For example, if the configured limit is 1000 TLS proxy sessions, and you purchase a 750-session UC license, then the first 250 IME sessions do not affect the sessions available for UC. If you need more than 250 sessions for IME, then the remaining 750 sessions of the platform limit are used on a first-come, first-served basis by UC and IME. • For a license part number ending in "K8", TLS proxy sessions are limited to 1000. • For a license part number ending in "K9", the TLS proxy limit depends on your configuration and the platform model. **Note** K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted. You might also use SRTP encryption sessions for your connections: • For a K8 license, SRTP sessions are limited to 250. • For a K9 license, there is no limit. **Note** Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.
Interfaces of all types, Max.	The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces.
IPS module	The IPS module license lets you run the IPS software module on the ASA. You must also purchase a separate IPS signature subscription; for failover, purchase a subscription for each unit. To obtain IPS signature support, you must purchase the ASA with IPS pre-installed (the part number must include "IPS"). The combined failover cluster license does not let you pair non-IPS and IPS units. For example, if you buy the IPS version of the ASA 5515-X (part number ASA5515-IPS-K9) and try to make a failover pair with a non-IPS version (part number ASA5515-K9), then Cisco will not let you obtain IPS signature updates for the ASA5515-K9 unit, even though it has an IPS module license inherited from the other unit.
Other VPN	Other VPN sessions include the following VPN types: • IPsec remote access VPN using IKEv1 • IPsec site-to-site VPN using IKEv1 • IPsec site-to-site VPN using IKEv2 This license is included in the Base license.

Table 5-16 License Notes (continued)

License	Notes
Total VPN (sessions), combined all types	• Although the maximum VPN sessions add up to more than the maximum VPN AnyConnect and Other VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the ASA, 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.

Table 5-16 License Notes (continued)

License	Notes
UC Phone Proxy sessions	The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit: • Phone Proxy • Presence Federation Proxy • Encrypted Voice Inspection Other applications that use TLS proxy sessions do not count toward the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license). Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used. You independently set the TLS proxy limit using the **tls-proxy maximum-sessions** command or in ASDM, using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. To view the limits of your model, enter the **tls-proxy maximum-sessions ?** command. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license. **Note** For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted. If you clear the configuration (using the **clear configure all** command, for example), then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the **tls-proxy maximum-sessions** command to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and enter the **write standby** command or in ASDM, use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning. You might also use SRTP encryption sessions for your connections: • For K8 licenses, SRTP sessions are limited to 250. • For K9 licenses, there is not limit. **Note** Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.
VPN Load Balancing	VPN load balancing requires a Strong Encryption (3DES/AES) License.

VPN License and Feature Compatibility

Table 5-17 shows how the VPN licenses and features can combine.

For a detailed list of the features supported by the AnyConnect Essentials license and AnyConnect Premium license, see *AnyConnect Secure Mobility Client Features, Licenses, and OSs*:

- Version 3.1: http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect31/feature/guide/anyconnect31features.html
- Version 3.0: http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect30/feature/guide/anyconnect30features.html
- Version 2.5: http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect25/feature/guide/anyconnect25features.html

Table 5-17 VPN License and Feature Compatibility

	Enable one of the following licenses:[1]	
Supported with:	**AnyConnect Essentials**	**AnyConnect Premium**
AnyConnect for Cisco VPN Phone	**No**	**Yes**
AnyConnect for Mobile[2]	**Yes**	**Yes**
Advanced Endpoint Assessment	**No**	**Yes**
AnyConnect Premium Shared	**No**	**Yes**
Client-based SSL VPN	**Yes**	**Yes**
Browser-based (clientless) SSL VPN	**No**	**Yes**
IPsec VPN	**Yes**	**Yes**
VPN Load Balancing	**Yes**	**Yes**
Cisco Secure Desktop	**No**	**Yes**

1. You can only have one license type active, either the AnyConnect Essentials license or the AnyConnect Premium license. By default, the ASA includes an AnyConnect Premium license for 2 sessions. If you install the AnyConnect Essentials license, then it is used by default. See the Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials pane to enable the Premium license instead.
2. Mobile Posture support is different for the AnyConnect Essentials vs. the AnyConnect Premium license. See Table 5-16 on page 5-18 for details.

Information About Feature Licenses

A license specifies the options that are enabled on a given ASA. 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 5-24
- Permanent License, page 5-24
- Time-Based Licenses, page 5-24

- Shared AnyConnect Premium Licenses, page 5-27
- Failover or ASA Cluster Licenses, page 5-31
- No Payload Encryption Models, page 5-33
- Licenses FAQ, page 5-33

Preinstalled License

By default, your ASA 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 "Monitoring Licenses" section on page 5-40 section to determine which licenses you have installed.

Permanent License

You can have one permanent activation key installed. The permanent activation key includes all licensed features in a single key. If you also install time-based licenses, the ASA combines the permanent and time-based licenses into a running license. See the "How Permanent and Time-Based Licenses Combine" section on page 5-25 for more information about how the ASA combines the licenses.

Time-Based Licenses

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

This section includes the following topics:

- Time-Based License Activation Guidelines, page 5-24
- How the Time-Based License Timer Works, page 5-25
- How Permanent and Time-Based Licenses Combine, page 5-25
- Stacking Time-Based Licenses, page 5-26
- Time-Based License Expiration, page 5-26

Time-Based License Activation Guidelines

- You can install multiple time-based licenses, including multiple licenses for the same feature. However, only one time-based license per feature can be *active* at a time. The inactive license remains installed, and ready for use. For example, if you install a 1000-session AnyConnect Premium license, and a 2500-session AnyConnect Premium license, then only one of these licenses can be active.
- If you activate an evaluation license that has multiple features in the key, then you cannot also activate another time-based license for one of the included features. For example, if an evaluation license includes the Botnet Traffic Filter and a 1000-session AnyConnect Premium license, you cannot also activate a standalone time-based 2500-session AnyConnect Premium license.

How the Time-Based License Timer Works

- The timer for the time-based license starts counting down when you activate it on the ASA.
- If you stop using the time-based license before it times out, then the timer halts. The timer only starts again when you reactivate the time-based license.
- If the time-based license is active, and you shut down the ASA, then the timer continues to count down. If you intend to leave the ASA in a shut down state for an extended period of time, then you should deactivate the time-based license before you shut down.

Note We suggest you do not change the system clock after you install the time-based license. If you set the clock to be a later date, then if you reload, the ASA 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 Permanent and Time-Based Licenses Combine

When you activate a time-based license, then features from both permanent and time-based licenses combine to form the running license. How the permanent and time-based licenses combine depends on the type of license. Table 5-18 lists the combination rules for each feature license.

Note Even when the permanent license is used, if the time-based license is active, it continues to count down.

Table 5-18 Time-Based License Combination Rules

Time-Based Feature	Combined License Rule
AnyConnect Premium Sessions	The higher value is used, either time-based or permanent. For example, if the permanent license is 1000 sessions, and the time-based license is 2500 sessions, then 2500 sessions are enabled. Typically, you will not install a time-based license that has less capability than the permanent license, but if you do so, then the permanent license is used.
Unified Communications Proxy Sessions	The time-based license sessions are added to the permanent sessions, up to the platform limit. For example, if the permanent license is 2500 sessions, and the time-based license is 1000 sessions, then 3500 sessions are enabled for as long as the time-based license is active.
Security Contexts	The time-based license contexts are added to the permanent contexts, up to the platform limit. For example, if the permanent license is 10 contexts, and the time-based license is 20 contexts, then 30 contexts are enabled for as long as the time-based license is active.
Botnet Traffic Filter	There is no permanent Botnet Traffic Filter license available; the time-based license is used.
All Others	The higher value is used, either time-based or permanent. For licenses that have a status of enabled or disabled, then the license with the enabled status is used. For licenses with numerical tiers, the higher value is used. Typically, you will not install a time-based license that has less capability than the permanent license, but if you do so, then the permanent license is used.

To view the combined license, see the "Monitoring Licenses" section on page 5-40.

Stacking Time-Based Licenses

In many cases, you might need to renew your time-based license and have a seamless transition from the old license to the new one. For features that are only available with a time-based license, it is especially important that the license not expire before you can apply the new license. The ASA allows you to *stack* time-based licenses so you do not have to worry about the license expiring or about losing time on your licenses because you installed the new one early.

When you install an identical time-based license as one already installed, then the licenses are combined, and the duration equals the combined duration.

For example:

1. You install a 52-week Botnet Traffic Filter license, and use the license for 25 weeks (27 weeks remain).
2. You then purchase another 52-week Botnet Traffic Filter license. When you install the second license, the licenses combine to have a duration of 79 weeks (52 weeks plus 27 weeks).

Similarly:

1. You install an 8-week 1000-session AnyConnect Premium license, and use it for 2 weeks (6 weeks remain).
2. You then install another 8-week 1000-session license, and the licenses combine to be 1000-sessions for 14 weeks (8 weeks plus 6 weeks).

If the licenses are not identical (for example, a 1000-session AnyConnect Premium license vs. a 2500-session license), then the licenses are *not* combined. Because only one time-based license per feature can be active, only one of the licenses can be active. See the "Activating or Deactivating Keys" section on page 5-37 for more information about activating licenses.

Although non-identical licenses do not combine, when the current license expires, the ASA automatically activates an installed license of the same feature if available. See the "Time-Based License Expiration" section on page 5-26 for more information.

Time-Based License Expiration

When the current license for a feature expires, the ASA automatically activates an installed license of the same feature if available. If there are no other time-based licenses available for the feature, then the permanent license is used.

If you have more than one additional time-based license installed for a feature, then the ASA uses the first license it finds; which license is used is not user-configurable and depends on internal operations. If you prefer to use a different time-based license than the one the ASA activated, then you must manually activate the license you prefer. See the "Activating or Deactivating Keys" section on page 5-37.

For example, you have a time-based 2500-session AnyConnect Premium license (active), a time-based 1000-session AnyConnect Premium license (inactive), and a permanent 500-session AnyConnect Premium license. While the 2500-session license expires, the ASA activates the 1000-session license. After the 1000-session license expires, the ASA uses the 500-session permanent license.

Shared AnyConnect Premium Licenses

A shared license lets you purchase a large number of AnyConnect Premium sessions and share the sessions as needed among a group of ASAs by configuring one of the ASAs 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 5-27
- Communication Issues Between Participant and Server, page 5-28
- Information About the Shared Licensing Backup Server, page 5-28
- Failover and Shared Licenses, page 5-29
- Maximum Number of Participants, page 5-30

Information About the Shared Licensing Server and Participants

The following steps describe how shared licenses operate:

1. Decide which ASA should be the shared licensing server, and purchase the shared licensing server license using that device serial number.
2. Decide which ASAs 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 ASA 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 ASA 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 ASA 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 5-29
- "Failover and Shared License Participants" section on page 5-30

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 ASA, 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.

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. 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 5-1).

Figure 5-1 Failover and Shared License Servers

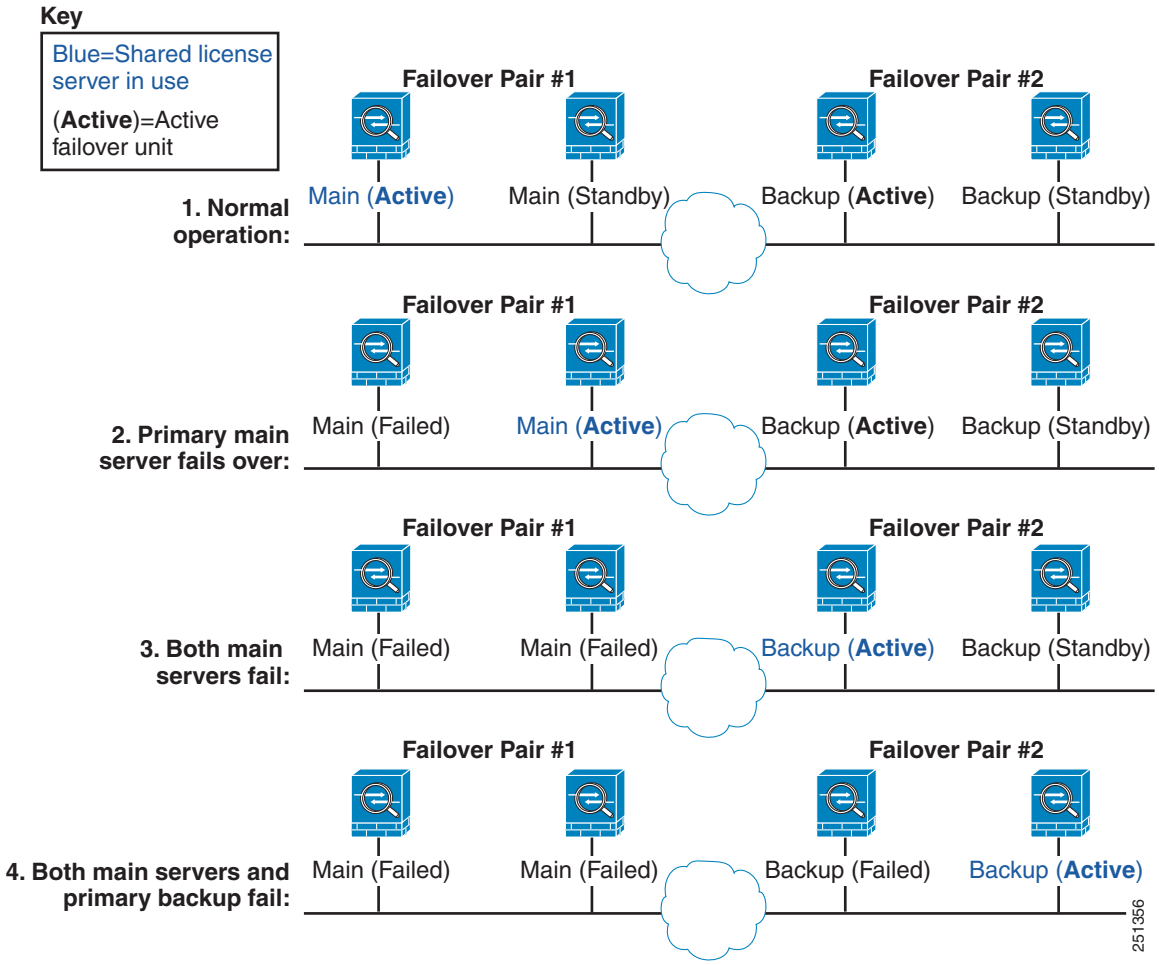

The standby backup server shares the same operating limits as the primary backup server; if the standby unit becomes active, it continues counting down where the primary unit left off. See the "Information About the Shared Licensing Backup Server" section on page 5-28 for more information.

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 ASA 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.

Failover or ASA Cluster Licenses

With some exceptions, failover and cluster units do not require the same license on each unit. For earlier versions, see the licensing document for your version.

This section includes the following topics:

- Failover License Requirements and Exceptions, page 5-31
- ASA Cluster License Requirements and Exceptions, page 5-31
- How Failover or ASA Cluster Licenses Combine, page 5-32
- Loss of Communication Between Failover or ASA Cluster Units, page 5-32
- Upgrading Failover Pairs, page 5-33

Failover License Requirements and Exceptions

Failover units do not require the same license on each unit.

Older versions of ASA software required that the licenses match on each unit. Starting with Version 8.3(1), you no longer need to install identical licenses. Typically, you buy a license only for the primary unit; for Active/Standby failover, the secondary unit inherits the primary license when it becomes active. If you have licenses on both units, they combine into a single running failover cluster license.

The exceptions to this rule include:

- Security Plus license for the ASA 5505, 5510, and 5512-X—The Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.
- Encryption license—Both units must have the same encryption license.
- IPS module license for the ASA 5512-X through ASA 5555-X—The IPS module license lets you run the IPS software module on the ASA. You must also purchase a separate IPS signature subscription for each unit. To obtain IPS signature support, you must purchase the ASA with IPS pre-installed (the part number must include “IPS”). The combined failover cluster license does not let you pair non-IPS and IPS units. For example, if you buy the IPS version of the ASA 5515-X (part number ASA5515-IPS-K9) and try to make a failover pair with a non-IPS version (part number ASA5515-K9), then you will not be able to obtain IPS signature updates for the ASA5515-K9 unit, even though it has an IPS module license inherited from the other unit.

Note A valid permanent key is required; in rare instances, your authentication key can be removed. If your key consists of all 0’s, then you need to reinstall a valid authentication key before failover can be enabled.

ASA Cluster License Requirements and Exceptions

Cluster units do not require the same license on each unit. Typically, you buy a license only for the master unit; slave units inherit the master license. If you have licenses on multiple units, they combine into a single running ASA cluster license.

The exceptions to this rule include:

- Clustering license—Each unit must have a clustering license.
- Encryption license—Each unit must have the same encryption license.

How Failover or ASA Cluster Licenses Combine

For failover pairs or ASA clusters, the licenses on each unit are combined into a single running cluster license. If you buy separate licenses for each unit, then the combined license uses the following rules:

- For licenses that have numerical tiers, such as the number of sessions, the values from each unit's licenses are combined up to the platform limit. If all licenses in use are time-based, then the licenses count down simultaneously.

  For example, for failover:

  – You have two ASAs with 10 AnyConnect Premium sessions installed on each; the licenses will be combined for a total of 20 AnyConnect Premium sessions.

  – You have two ASA 5520 ASAs with 500 AnyConnect Premium sessions each; because the platform limit is 750, the combined license allows 750 AnyConnect Premium sessions.

  Note In the above example, if the AnyConnect Premium licenses are time-based, you might want to disable one of the licenses so you do not "waste" a 500 session license from which you can only use 250 sessions because of the platform limit.

  – You have two ASA 5540 ASAs, one with 20 contexts and the other with 10 contexts; the combined license allows 30 contexts. For Active/Active failover, the contexts are divided between the two units. One unit can use 18 contexts and the other unit can use 12 contexts, for example, for a total of 30.

  For example, for ASA clustering:

  – You have three units with 50 contexts each, and one unit with the default 2 contexts. The combined license allows 152 contexts to be divided up amongst the four cluster members. Therefore, you can configure up to 38 contexts on the master unit; each slave unit will also have 38 contexts through configuration replication.

- For licenses that have a status of enabled or disabled, then the license with the enabled status is used.
- For time-based licenses that are enabled or disabled (and do not have numerical tiers), the duration is the combined duration of all licenses. The primary/master unit counts down its license first, and when it expires, the secondary/slave unit(s) start counting down its license, and so on. This rule also applies to Active/Active failover and ASA clustering, even though all units are actively operating.

  For example, if you have 48 weeks left on the Botnet Traffic Filter license on two units, then the combined duration is 96 weeks.

To view the combined license, see the "Monitoring Licenses" section on page 5-40.

Loss of Communication Between Failover or ASA Cluster Units

If the units lose communication for more than 30 days, then each unit reverts to the license installed locally. During the 30-day grace period, the combined running license continues to be used by all units.

If you restore communication during the 30-day grace period, then for time-based licenses, the time elapsed is subtracted from the primary/master license; if the primary/master license becomes expired, only then does the secondary/slave license start to count down.

If you do not restore communication during the 30-day period, then for time-based licenses, time is subtracted from all unit licenses, if installed. They are treated as separate licenses and do not benefit from the combined license. The time elapsed includes the 30-day grace period.

For example:

1. You have a 52-week Botnet Traffic Filter license installed on two units. The combined running license allows a total duration of 104 weeks.
2. The units operate as a failover unit/ASA cluster for 10 weeks, leaving 94 weeks on the combined license (42 weeks on the primary/master, and 52 weeks on the secondary/slave).
3. If the units lose communication (for example the primary/master unit fails), the secondary/slave unit continues to use the combined license, and continues to count down from 94 weeks.
4. The time-based license behavior depends on when communication is restored:

- Within 30 days—The time elapsed is subtracted from the primary/master unit license. In this case, communication is restored after 4 weeks. Therefore, 4 weeks are subtracted from the primary/master license leaving 90 weeks combined (38 weeks on the primary, and 52 weeks on the secondary).
- After 30 days—The time elapsed is subtracted from both units. In this case, communication is restored after 6 weeks. Therefore, 6 weeks are subtracted from both the primary/master and secondary/slave licenses, leaving 84 weeks combined (36 weeks on the primary/master, and 46 weeks on the secondary/slave).

Upgrading Failover Pairs

Because failover pairs do not require the same license on both units, you can apply new licenses to each unit without any downtime. If you apply a permanent license that requires a reload (see Table 5-19 on page 5-37), then you can fail over to the other unit while you reload. If both units require reloading, then you can reload them separately so you have no downtime.

No Payload Encryption Models

You can purchase some models with No Payload Encryption. For export to some countries, payload encryption cannot be enabled on the Cisco ASA series. The ASA software senses a No Payload Encryption model, and disables the following features:

- Unified Communications
- VPN

You can still install the Strong Encryption (3DES/AES) license for use with management connections. For example, you can use ASDM HTTPS/SSL, SSHv2, Telnet and SNMPv3. You can also download the dynamic database for the Botnet Traffic Filer (which uses SSL).

When you view the license (see the "Monitoring Licenses" section on page 5-40), VPN and Unified Communications licenses will not be listed.

Licenses FAQ

Q. Can I activate multiple time-based licenses, for example, AnyConnect Premium and Botnet Traffic Filter?

A. Yes. You can use one time-based license per feature at a time.

Q. Can I "stack" time-based licenses so that when the time limit runs out, it will automatically use the next license?

A. Yes. For identical licenses, the time limit is combined when you install multiple time-based licenses. For non-identical licenses (for example, a 1000-session AnyConnect Premium license and a 2500-session license), the ASA automatically activates the next time-based license it finds for the feature.

Q. Can I install a new permanent license while maintaining an active time-based license?

A. Yes. Activating a permanent license does not affect time-based licenses.

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 has the same running license as the primary unit; in the case of the shared licensing server, they require a server license. The backup server requires a participant license. The backup server 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?

A. No. Starting with Version 8.3(1), you do not have to have matching licenses on both units. Typically, you buy a license only for the primary unit; the secondary unit inherits the primary license when it becomes active. In the case where you also have a separate license on the secondary unit (for example, if you purchased matching licenses for pre-8.3 software), the licenses are combined into a running failover cluster license, up to the model limits.

Q. Can I use a time-based or permanent AnyConnect Premium license in addition to a shared AnyConnect Premium license?

A. Yes. The shared license is used only after the sessions from the locally installed license (time-based or permanent) are used up. **Note**: On the shared licensing server, the permanent AnyConnect Premium license is not used; you can however use a time-based license at the same time as the shared licensing server license. In this case, the time-based license sessions are available for local AnyConnect Premium 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

- Shared licenses are not supported in Active/Active mode. See the "Failover and Shared Licenses" section on page 5-29 for more information.
- Failover units do not require the same license on each unit.

Older versions of ASA software required that the licenses match on each unit. Starting with Version 8.3(1), you no longer need to install identical licenses. Typically, you buy a license only for the primary unit; for Active/Standby failover, the secondary unit inherits the primary license when it becomes active. If you have licenses on both units, they combine into a single running failover cluster license.

Note Failover units do require the same RAM on both units.

- For the ASA 5505 and 5510, both units require the Security Plus license; the Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.

Upgrade and Downgrade Guidelines

Your activation key remains compatible if you upgrade to the latest version from any previous version. However, you might have issues if you want to maintain downgrade capability:

- Downgrading to Version 8.1 or earlier—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 ASA 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.
- Downgrading to Version 8.2 or earlier—Version 8.3 introduced more robust time-based key usage as well as failover license changes:
  - If you have more than one time-based activation key active, when you downgrade, only the most recently activated time-based key can be active. Any other keys are made inactive. If the last time-based license is for a feature introduced in 8.3, then that license still remains the active license even though it cannot be used in earlier versions. Reenter the permanent key or a valid time-based key.
  - If you have mismatched licenses on a failover pair, then downgrading will disable failover. Even if the keys are matching, the license used will no longer be a combined license.
  - If you have one time-based license installed, but it is for a feature introduced in 8.3, then after you downgrade, that time-based license remains active. You need to reenter the permanent key to disable the time-based license.

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, and it is covered by Cisco TAC, 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.

- On a single unit, 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. (You may be able to purchase a larger license at an upgrade price, for example from 25 sessions to 75 sessions; this kind of upgrade should be distinguished from adding two separate licenses together).
- Although you can activate all license types, some features are incompatible with each other. In the case of the AnyConnect Essentials license, the license is incompatible with the following licenses: AnyConnect Premium license, shared AnyConnect Premium 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 Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials pane.

Configuring Licenses

This section includes the following topics:

- Obtaining an Activation Key, page 5-36
- Activating or Deactivating Keys, page 5-37
- Configuring a Shared License, page 5-38

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 AnyConnect Premium sessions.

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

Detailed Steps

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

Step 2 If you are not already registered with Cisco.com, create an account.

Step 3 Go to the following licensing website:

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

Step 4 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 ASA
- Your e-mail address

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

Step 5 If you have additional Product Authorization Keys, repeat Step 4 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.

Activating or Deactivating Keys

This section describes how to enter a new activation key, and how to activate and deactivate time-based keys.

Prerequisites

- If you are already in multiple context mode, enter the activation key in the system execution space.
- Some permanent licenses require you to reload the ASA after you activate them. Table 5-19 lists the licenses that require reloading.

Table 5-19 Permanent License Reloading Requirements

Model	License Action Requiring Reload
ASA 5505, 5510	Changing between the Base and Security Plus license.
All models	Changing the Encryption license.
All models	Downgrading any permanent license (for example, going from 10 contexts to 2 contexts).

Limitations and Restrictions

Your activation key remains compatible if you upgrade to the latest version from any previous version. However, you might have issues if you want to maintain downgrade capability:

- Downgrading to Version 8.1 or earlier—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 ASA 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.
- Downgrading to Version 8.2 or earlier—Version 8.3 introduced more robust time-based key usage as well as failover license changes:
  - If you have more than one time-based activation key active, when you downgrade, only the most recently activated time-based key can be active. Any other keys are made inactive.
  - If you have mismatched licenses on a failover pair, then downgrading will disable failover. Even if the keys are matching, the license used will no longer be a combined license.

Detailed Steps

Step 1 Choose **Configuration > Device Management**, and then choose the **Licensing > Activation Key** or **Licensing Activation Key** pane, depending on your model.

Step 2 To enter a new activation key, either permanent or time-based, 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 install one permanent key, and multiple time-based keys. If you enter a new permanent key, it overwrites the already installed one. If you enter a new time-based key, then it is active by default and displays in the Time-based License Keys Installed table. The last time-based key that you activate for a given feature is the active one.

Step 3 To activate or deactivate an installed time-based key, choose the key in the Time-based License Keys Installed table, and click either **Activate** or **Deactivate**.

You can only have one time-based key active for each feature. See the "Time-Based Licenses" section on page 5-24 for more information.

Step 4 Click **Update Activation Key**.

Some permanent licenses require you to reload the ASA after entering the new activation key. See Table 5-19 on page 5-37 for a list of licenses that need reloading. You will be prompted to reload if it is required.

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 AnyConnect Premium Licenses" section on page 5-27.

This section includes the following topics:

- Configuring the Shared Licensing Server, page 5-38
- Configuring the Shared Licensing Participant and the Optional Backup Server, page 5-39

Configuring the Shared Licensing Server

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

Prerequisites

The server must have a shared licensing server 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.

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 5-39.

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 Licenses

This section includes the following topics:

- Viewing Your Current License, page 5-40
- Monitoring the Shared License, page 5-41

Viewing Your Current License

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

Guidelines

If you have a No Payload Encryption model, then you view the license, VPN and Unified Communications licenses will not be listed. See the "No Payload Encryption Models" section on page 5-33 for more information.

Detailed Steps

Step 1 To view the running license, which is a combination of the permanent license and any active time-based licenses, choose the **Configuration > Device Management > Licensing > Activation Key** pane and view the Running Licenses area.

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

For a failover pair, the running license shown is the combined license from the primary and secondary units. See the "How Failover or ASA Cluster Licenses Combine" section on page 5-32 for more information. For time-based licenses with numerical values (the duration is not combined), the License Duration column displays the shortest time-based license from either the primary or secondary unit; when that license expires, the license duration from the other unit displays.

Step 2 (Optional) To view time-based license details, such as the features included in the license and the duration, in the Time-Based License Keys Installed area, choose a license key, and then click **Show License Details**.

Step 3 (Optional) For a failover unit, to view the license installed on this unit (and not the combined license from both primary and secondary units), in the Running Licenses area, click **Show information of license specifically purchased for this device alone**.

Monitoring the Shared License

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

Feature History for Licensing

Table 5-20 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 5-20 Feature History for Licensing

Feature Name	Platform 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 interfaces for the Base license on the ASA 5510	7.2(2)	For the Base license on the ASA 5510, the maximum number of interfaces was increased from 3 plus a management interface to unlimited interfaces.
Increased VLANs	7.2(2)	The maximum number of VLANs for the Security Plus license on the ASA 5505 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 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).

Table 5-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 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 ASA. The ASA 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 was introduced. It lets Windows mobile devices connect to the ASA using the AnyConnect client.
Time-based Licenses	8.0(4)/8.1(2)	Support for time-based licenses was introduced.
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. Phone Proxy, Presence Federation Proxy, and Encrypted Voice Inspection applications use TLS proxy sessions for their connections. Each TLS proxy session is counted against the UC license limit. All of these applications are licensed under the UC Proxy umbrella, and can be mixed and matched. This feature is not available in Version 8.1.

Table 5-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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.
AnyConnect Essentials License	8.2(1)	The AnyConnect Essentials License was introduced. This license enables AnyConnect VPN client access to the ASA. This license does not support browser-based SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium 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 license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given ASA: AnyConnect Premium license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium licenses on different ASAs in the same network. By default, the ASA uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials pane.
SSL VPN license changed to AnyConnect Premium SSL VPN Edition license	8.2(1)	The SSL VPN license name was changed to the AnyConnect Premium SSL VPN Edition license.
Shared Licenses for SSL VPN	8.2(1)	Shared licenses for SSL VPN were introduced. Multiple ASAs can share a pool of SSL VPN sessions on an as-needed basis.
Mobility Proxy application no longer requires Unified Communications Proxy license	8.2(2)	The Mobility Proxy no longer requires the UC Proxy license.
10 GE I/O license for the ASA 5585-X with SSP-20	8.2(3)	We introduced the 10 GE I/O license for the ASA 5585-X with SSP-20 to enable 10-Gigabit Ethernet speeds for the fiber ports. The SSP-60 supports 10-Gigabit Ethernet speeds by default. **Note** The ASA 5585-X is not supported in 8.3(x).
10 GE I/O license for the ASA 5585-X with SSP-10	8.2(4)	We introduced the 10 GE I/O license for the ASA 5585-X with SSP-10 to enable 10-Gigabit Ethernet speeds for the fiber ports. The SSP-40 supports 10-Gigabit Ethernet speeds by default. **Note** The ASA 5585-X is not supported in 8.3(x).

Table 5-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
Non-identical failover licenses	8.3(1)	Failover licenses no longer need to be identical on each unit. The license used for both units is the combined license from the primary and secondary units. We modified the following screen: Configuration > Device Management > Licensing > Activation Key.
Stackable time-based licenses	8.3(1)	Time-based licenses are now stackable. In many cases, you might need to renew your time-based license and have a seamless transition from the old license to the new one. For features that are only available with a time-based license, it is especially important that the license not expire before you can apply the new license. The ASA allows you to *stack* time-based licenses so you do not have to worry about the license expiring or about losing time on your licenses because you installed the new one early.
Intercompany Media Engine License	8.3(1)	The IME license was introduced.
Multiple time-based licenses active at the same time	8.3(1)	You can now install multiple time-based licenses, and have one license per feature active at a time. The following screen was modified: Configuration > Device Management > Licensing > Activation Key.
Discrete activation and deactivation of time-based licenses.	8.3(1)	You can now activate or deactivate time-based licenses using a command. The following screen was modified: Configuration > Device Management > Licensing > Activation Key.
AnyConnect Premium SSL VPN Edition license changed to AnyConnect Premium SSL VPN license	8.3(1)	The AnyConnect Premium SSL VPN Edition license name was changed to the AnyConnect Premium SSL VPN license.
No Payload Encryption image for export	8.3(2)	If you install the No Payload Encryption software on the ASA 5505 through 5550, then you disable Unified Communications, strong encryption VPN, and strong encryption management protocols. **Note** This special image is only supported in 8.3(x); for No Payload Encryption support in 8.4(1) and later, you need to purchase a special hardware version of the ASA.
Increased contexts for the ASA 5550, 5580, and 5585-X	8.4(1)	For the ASA 5550 and ASA 5585-X with SSP-10, the maximum contexts was increased from 50 to 100. For the ASA 5580 and 5585-X with SSP-20 and higher, the maximum was increased from 50 to 250.
Increased VLANs for the ASA 5580 and 5585-X	8.4(1)	For the ASA 5580 and 5585-X, the maximum VLANs was increased from 250 to 1024.

Table 5-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
Increased connections for the ASA 5580 and 5585-X	8.4(1)	We increased the firewall connection limits: • ASA 5580-20—1,000,000 to 2,000,000. • ASA 5580-40—2,000,000 to 4,000,000. • ASA 5585-X with SSP-10: 750,000 to 1,000,000. • ASA 5585-X with SSP-20: 1,000,000 to 2,000,000. • ASA 5585-X with SSP-40: 2,000,000 to 4,000,000. • ASA 5585-X with SSP-60: 2,000,000 to 10,000,000.
AnyConnect Premium SSL VPN license changed to AnyConnect Premium license	8.4(1)	The AnyConnect Premium SSL VPN license name was changed to the AnyConnect Premium license. The license information display was changed from "SSL VPN Peers" to "AnyConnect Premium Peers."
Increased AnyConnect VPN sessions for the ASA 5580	8.4(1)	The AnyConnect VPN session limit was increased from 5,000 to 10,000.
Increased Other VPN sessions for the ASA 5580	8.4(1)	The other VPN session limit was increased from 5,000 to 10,000.
IPsec remote access VPN using IKEv2	8.4(1)	IPsec remote access VPN using IKEv2 was added to the AnyConnect Essentials and AnyConnect Premium licenses. IKEv2 site-to-site sessions were added to the Other VPN license (formerly IPsec VPN). The Other VPN license is included in the Base license.
No Payload Encryption hardware for export	8.4(1)	For models available with No Payload Encryption (for example, the ASA 5585-X), the ASA software disables Unified Communications and VPN features, making the ASA available for export to certain countries.
Dual SSPs for SSP-20 and SSP-40	8.4(2)	For SSP-40 and SSP-60, you can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired. When using two SSPs in the chassis, VPN is not supported; note, however, that VPN has not been disabled.
IPS Module license for the ASA 5512-X through ASA 5555-X	8.6(1)	The IPS SSP software module on the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X requires the IPS module license.
Clustering license for the ASA 5580 and ASA 5585-X.	9.0(1)	A clustering license was added for the ASA 5580 and ASA 5585-X.
Support for VPN on the ASASM	9.0(1)	The ASASM now supports all VPN features.
Unified communications support on the ASASM	9.0(1)	The ASASM now supports all Unified Communications features.

Configuring the Transparent or Routed Firewall

This chapter describes how to set the firewall mode to routed or transparent, as well as how the firewall works in each firewall mode. This chapter also includes information about customizing the transparent firewall operation.

You can set the firewall mode independently for each context in multiple context mode.

This chapter includes the following sections:

- Configuring the Firewall Mode, page 6-1
- Configuring ARP Inspection for the Transparent Firewall, page 6-10
- Customizing the MAC Address Table for the Transparent Firewall, page 6-14
- Firewall Mode Examples, page 6-17

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, page 6-1
- Licensing Requirements for the Firewall Mode, page 6-6
- Default Settings, page 6-6
- Guidelines and Limitations, page 6-6
- Setting the Firewall Mode (Single Mode), page 6-8
- Feature History for Firewall Mode, page 6-9

Information About the Firewall Mode

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

- Information About Routed Firewall Mode, page 6-2
- Information About Transparent Firewall Mode, page 6-2

Information About Routed Firewall Mode

In routed mode, the ASA is considered to be a router hop in the network. Routed mode supports many interfaces. Each interface is on a different subnet. You can share interfaces between contexts.

The ASA acts as a router between connected networks, and each interface requires an IP address on a different subnet. The ASA supports multiple dynamic routing protocols. However, we recommend using the advanced routing capabilities of the upstream and downstream routers instead of relying on the ASA 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
- Bridge Groups, page 6-2
- Management Interface (ASA 5510 and Higher), page 6-3
- Allowing Layer 3 Traffic, page 6-3
- Allowed MAC Addresses, page 6-3
- Passing Traffic Not Allowed in Routed Mode, page 6-3
- BPDU Handling, page 6-4
- MAC Address vs. Route Lookups, page 6-4
- Using the Transparent Firewall in Your Network, page 6-5

Transparent Firewall Network

The ASA connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

Bridge Groups

If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the ASA, and traffic must exit the ASA before it is routed by an external router back to another bridge group in the ASA. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration. For complete security policy separation, use security contexts with one bridge group in each context.

Note Each bridge group requires a management IP address. The ASA uses this IP address as the source address for packets originating from the bridge group. The management IP address must be on the same subnet as the connected network. For another method of management, see the "Management Interface (ASA 5510 and Higher)" section on page 6-3.

The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.

Management Interface (ASA 5510 and Higher)

In addition to each bridge group management IP address, you can add a separate Management *slot*/*port* interface that is not part of any bridge group, and that allows only management traffic to the ASA. For more information, see the "Management Interface" section on page 15-2.

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 51, "Configuring Access Rules," 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 ASA 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).

Non-IP traffic (for example AppleTalk, IPX, BPDUs, and MPLS) can be configured to go through using an EtherType access list.

Note The transparent mode ASA 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.

Passing Traffic For Routed-Mode Features

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. You can also 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 ASA.

BPDU Handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default. To block BPDUs, you need to configure an EtherType access list to deny them. If you are using failover, you might want to block BPDUs to prevent the switch port from going into a blocking state when the topology changes. See the "Transparent Firewall Mode Requirements" section on page 11-15 for more information.

MAC Address vs. Route Lookups

When the ASA runs in transparent mode, the outgoing interface of a packet is determined by performing a MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for the following traffic types:

- Traffic originating on the ASA—For example, if your syslog server is located on a remote network, you must use a static route so the ASA can reach that subnet.
- Traffic that is at least one hop away from the ASA with NAT enabled—The ASA needs to perform a route lookup; you need to add a static route on the ASA for the real host address.
- Voice over IP (VoIP) traffic with inspection enabled, and the endpoint is at least one hop away from the ASA—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 ASA for the H.323 gateway for successful call completion.
- VoIP or DNS traffic with inspection enabled, with NAT enabled, and the embedded address is at least one hop away from the ASA—To successfully translate the IP address inside VoIP and DNS packets, the ASA needs to perform a route lookup; you need to add a static route on the ASA for the real host address that is embedded in the packet.

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

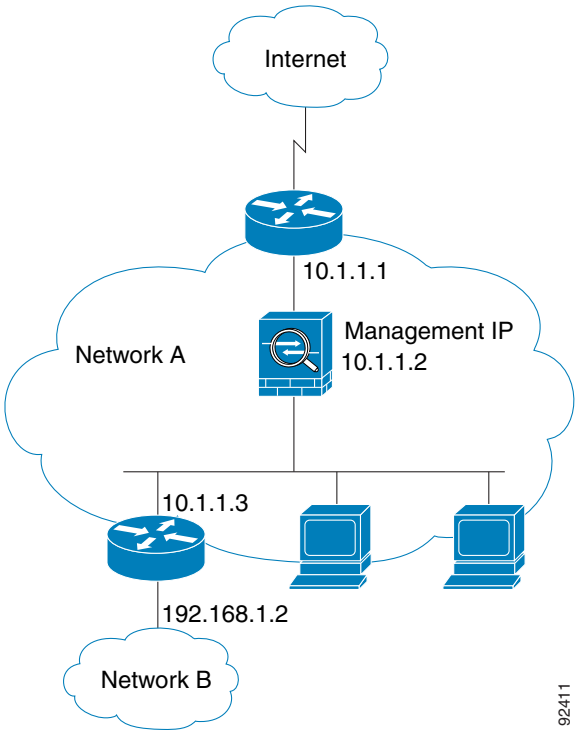

Figure 6-2 shows two networks connected to the ASA, which has two bridge groups.

Figure 6-2 Transparent Firewall Network with Two Bridge Groups

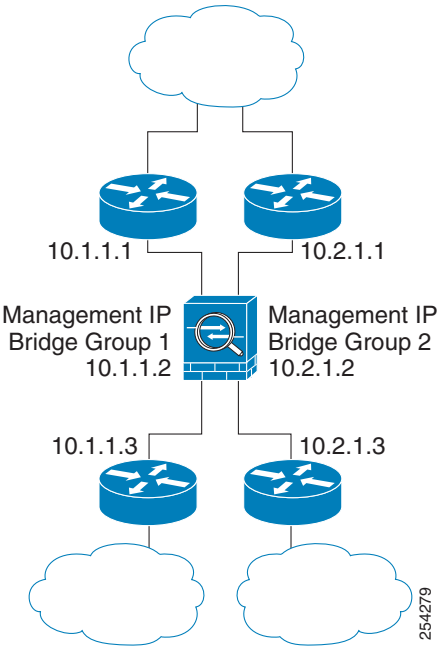

Licensing Requirements for the Firewall Mode

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

Set the firewall mode per context.

Transparent Firewall Guidelines

Follow these guidelines when planning your transparent firewall network:

- In transparent firewall mode, the management interface updates the MAC address table in the same manner as a data interface; therefore you should not connect both a management and a data interface to the same switch unless you configure one of the switch ports as a routed port (by default Cisco Catalyst switches share a MAC address for all VLAN switch ports). Otherwise, if traffic arrives on the management interface from the physically-connected switch, then the ASA updates the MAC address table to use the *management* interface to access the switch, instead of the data interface. This action causes a temporary traffic interruption; the ASA will not re-update the MAC address table for packets from the switch to the data interface for at least 30 seconds for security reasons.
- Each directly-connected network must be on the same subnet.
- Do not specify the bridge group management IP address as the default gateway for connected devices; devices need to specify the router on the other side of the ASA as the default gateway.
- The default route for the transparent firewall, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a static route that identifies the network from which you expect management traffic.

See the "Guidelines and Limitations" section on page 18-5 for more guidelines.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

- When you change firewall modes, the ASA 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 (Single Mode)" section on page 6-8 for information about backing up your configuration file.
- If you download a text configuration to the ASA that changes the mode with the **firewall transparent** command, be sure to put the command at the top of the configuration; the ASA 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 ASA 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 ASA. You can also allow dynamic routing protocols through the ASA using an extended access list.
Multicast IP routing	You can allow multicast traffic through the ASA 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 ASA. You can pass VPN traffic through the ASA using an extended access list, but it does not terminate non-management connections. SSL VPN is also not supported.
Unified Communications	—

Setting the Firewall Mode (Single Mode)

This section describes how to change the firewall mode using the CLI. For single mode, you cannot change the mode in ASDM. For multiple mode, see the "Configuring a Security Context" section on page 9-20.

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 ASA clears the running configuration (see the "Guidelines and Limitations" section on page 6-6 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 ASA using the console port in any case.
- Set the mode within the context.

Detailed Steps

Command	Purpose

`firewall transparent`
 	
Example:
`hostname(config)# firewall transparent`
	Sets the firewall mode to transparent. To change the mode to routed, enter the **no firewall transparent** command.
	
	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. We introduced the following commands: **firewall transparent**, **show firewall**. You cannot set the firewall mode in ASDM; you must use the command-line interface.

Table 6-2 Feature History for Firewall Mode (continued)

Feature Name	Releases	Feature Information
Transparent firewall bridge groups	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to 8 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. **Note** Although you can configure multiple bridge groups on the ASA 5505, the restriction of 2 data interfaces in transparent mode on the ASA 5505 means you can only effectively use 1 bridge group. We modified or introduced the following screens: Configuration > Device Setup > Interfaces Configuration > Device Setup > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interfaces > Add/Edit Interface
Mixed firewall mode support in multiple context mode	8.5(1)/9.0(1)	You can set the firewall mode independently for each security context in multiple context mode, so some can run in transparent mode while others run in routed mode. We modified the following command: **firewall transparent**. For single mode, you cannot set the firewall mode in ASDM; you must use the command-line interface. For multiple mode, we modified the following screen: Configuration > Context Management > Security Contexts.

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, page 6-11
- Licensing Requirements for ARP Inspection, page 6-11
- Default Settings, page 6-11
- Guidelines and Limitations, page 6-11
- Configuring ARP Inspection, page 6-12
- Feature History for ARP Inspection, page 6-14

Information About ARP Inspection

By default, all ARP packets are allowed through the ASA. You can control the flow of ARP packets by enabling ARP inspection.

When you enable ARP inspection, the ASA 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 ASA drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the ASA 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 ASA.

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 6-12
- Adding a Static ARP Entry, page 6-12
- Enabling ARP Inspection, page 6-13

Task Flow for Configuring ARP Inspection

To configure ARP Inspection, perform the following steps:

Step 1 Add static ARP entries according to the "Adding a Static ARP Entry" section on page 6-12. 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 6-13.

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 ASA, such as management traffic.

Detailed Steps

Step 1 Choose the **Configuration > Device Management > Advanced > 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 ASA 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 (Optional; 8.4(5) only) To allow non-connected subnets, check the **Allow non-connected subnets** check box. The ASA ARP cache only contains entries from directly-connected subnets by default. You can enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries.

You may want to use this feature if you use:

- Secondary subnets.
- Proxy ARP on adjacent routes for traffic forwarding.

Step 4 Click **Add**.

The Add ARP Static Configuration dialog box appears.

Step 5 From the Interface drop-down list, choose the interface attached to the host network.

Step 6 In the IP Address field, enter the IP address of the host.

Step 7 In the MAC Address field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 8 To perform proxy ARP for this address, check the **Proxy ARP** check box.

If the ASA receives an ARP request for the specified IP address, then it responds with the specified MAC address.

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

What to Do Next

Enable ARP inspection according to the "Enabling ARP Inspection" section on page 6-13.

Enabling ARP Inspection

This section describes how to enable ARP inspection.

Detailed Steps

Step 1 Choose the **Configuration > Device Management > Advanced > 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 ASA drops the packet.

If you uncheck this check box, all non-matching packets are dropped, which restricts ARP through the ASA 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. We introduced the following commands: **arp**, **arp-inspection**, and **show arp-inspection**.
ARP cache additions for non-connected subnets	8.4(5)	The ASA ARP cache only contains entries from directly-connected subnets by default. You can now enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries. You may want to use this feature if you use: • Secondary subnets. • Proxy ARP on adjacent routes for traffic forwarding. We modified the following screen: Configuration > Device Management > Advanced > ARP > ARP Static Table. *This feature is not available in 8.5(1), 8.6(1), or 9.0(1).*

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 6-15
- Licensing Requirements for the MAC Address Table, page 6-15
- Default Settings, page 6-15
- Guidelines and Limitations, page 6-15
- Configuring the MAC Address Table, page 6-16
- Feature History for the MAC Address Table, page 6-17

Information About the MAC Address Table

The ASA 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 ASA, the ASA adds the MAC address to its table. The table associates the MAC address with the source interface so that the ASA knows to send any packets addressed to the device out the correct interface.

The ASA 5505 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 only discusses the *bridge* MAC address table, which maintains the MAC address-to-VLAN interface mapping for traffic that passes between VLANs.

Because the ASA is a firewall, if the destination MAC address of a packet is not in the table, the ASA 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 ASA generates an ARP request for the destination IP address, so that the ASA can learn which interface receives the ARP response.
- Packets for remote devices—The ASA generates a ping to the destination IP address so that the ASA can learn which interface receives the ping reply.

The original packet is dropped.

Licensing Requirements for the MAC Address Table

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 ASA 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.

Additional Guidelines

In transparent firewall mode, the management interface updates the MAC address table in the same manner as a data interface; therefore you should not connect both a management and a data interface to the same switch unless you configure one of the switch ports as a routed port (by default Cisco Catalyst switches share a MAC address for all VLAN switch ports). Otherwise, if traffic arrives on the management interface from the physically-connected switch, then the ASA updates the MAC address table to use the *management* interface to access the switch, instead of the data interface. This action causes a temporary traffic interruption; the ASA will not re-update the MAC address table for packets from the switch to the data interface for at least 30 seconds for security reasons.

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 6-16
- Disabling MAC Address Learning, page 6-16

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 ASA 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 6-12), 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 ASA 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 ASA.

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. We introduced the following commands: **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 ASA and includes the following topics:

- How Data Moves Through the ASA in Routed Firewall Mode, page 6-17
- How Data Moves Through the Transparent Firewall, page 6-23

How Data Moves Through the ASA in Routed Firewall Mode

This section describes how data moves through the ASA in routed firewall mode and includes the following topics:

- An Inside User Visits a Web Server, page 6-18
- An Outside User Visits a Web Server on the DMZ, page 6-19
- An Inside User Visits a Web Server on the DMZ, page 6-20
- An Outside User Attempts to Access an Inside Host, page 6-20
- A DMZ User Attempts to Access an Inside Host, page 6-22

An Inside User Visits a Web Server

Figure 6-3 shows an inside user accessing an outside web server.

Figure 6-3 Inside to Outside

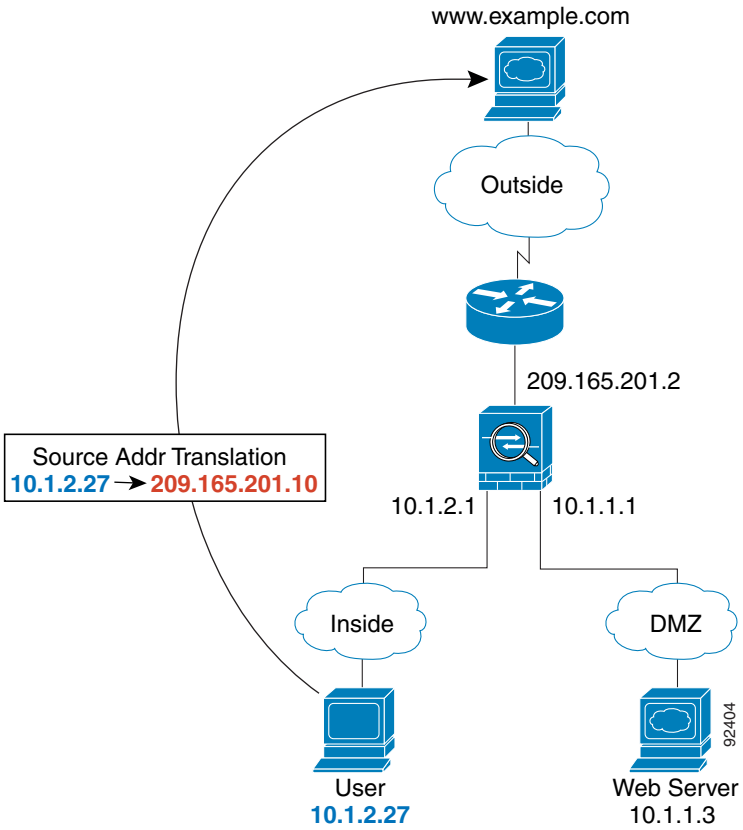

The following steps describe how data moves through the ASA (see Figure 6-3):

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA 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 ASA 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 ASA, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by untranslating the global destination address to the local user address, 10.1.2.27.
6. The ASA forwards the packet to the inside user.

An Outside User Visits a Web Server on the DMZ

Figure 6-4 shows an outside user accessing the DMZ web server.

Figure 6-4 Outside to DMZ

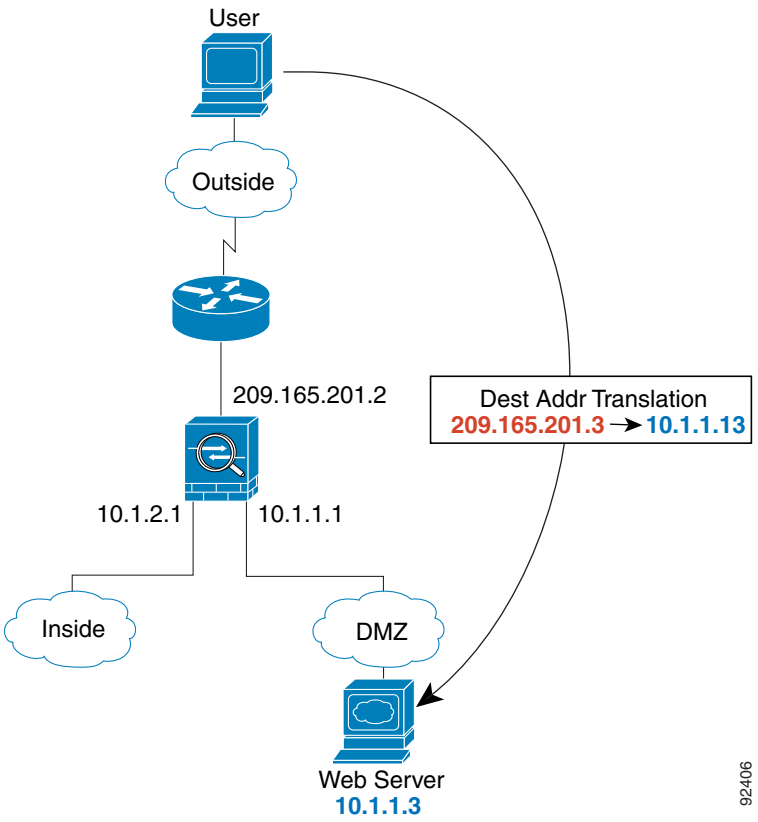

The following steps describe how data moves through the ASA (see Figure 6-4):

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 ASA receives the packet and untranslates the destination address to the local address 10.1.1.3.
3. Because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the ASA first classifies the packet to a context.
4. The ASA 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 ASA and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by translating the local source address to 209.165.201.3.
6. The ASA forwards the packet to the outside user.

An Inside User Visits a Web Server on the DMZ

Figure 6-5 shows an inside user accessing the DMZ web server.

Figure 6-5 Inside to DMZ

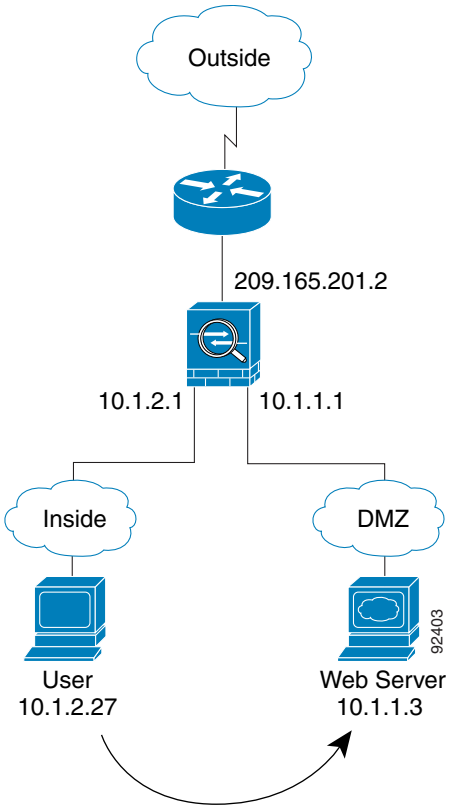

The following steps describe how data moves through the ASA (see Figure 6-5):

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 ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA 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 ASA forwards the packet to the inside user.

An Outside User Attempts to Access an Inside Host

Figure 6-6 shows an outside user attempting to access the inside network.

Figure 6-6 Outside to Inside

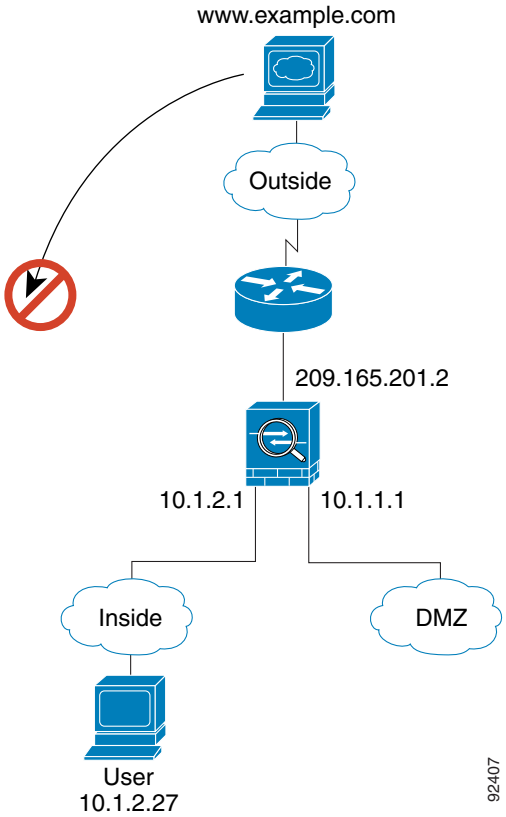

The following steps describe how data moves through the ASA (see Figure 6-6):

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 ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according to the security policy (access lists, filters, AAA).

3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

   If the outside user is attempting to attack the inside network, the ASA 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 6-7 shows a user in the DMZ attempting to access the inside network.

Figure 6-7 DMZ to Inside

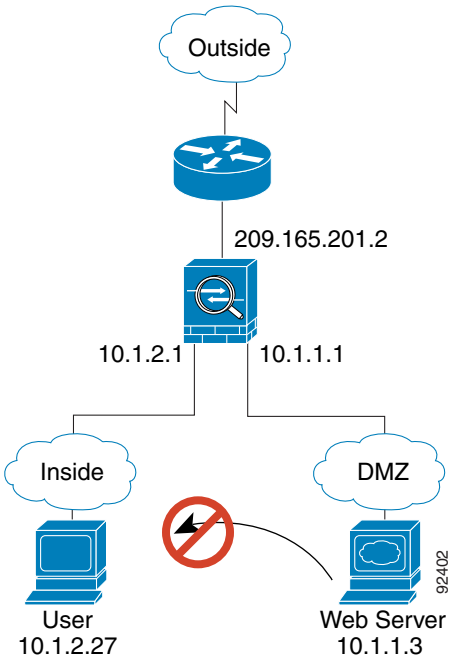

The following steps describe how data moves through the ASA (see Figure 6-7):

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 ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according to the security policy (access lists, filters, AAA).

   The packet is denied, and the ASA drops the packet and logs the connection attempt.

How Data Moves Through the Transparent Firewall

Figure 6-8 shows a typical transparent firewall implementation with an inside network that contains a public web server. The ASA 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 6-8 Typical Transparent Firewall Data Path

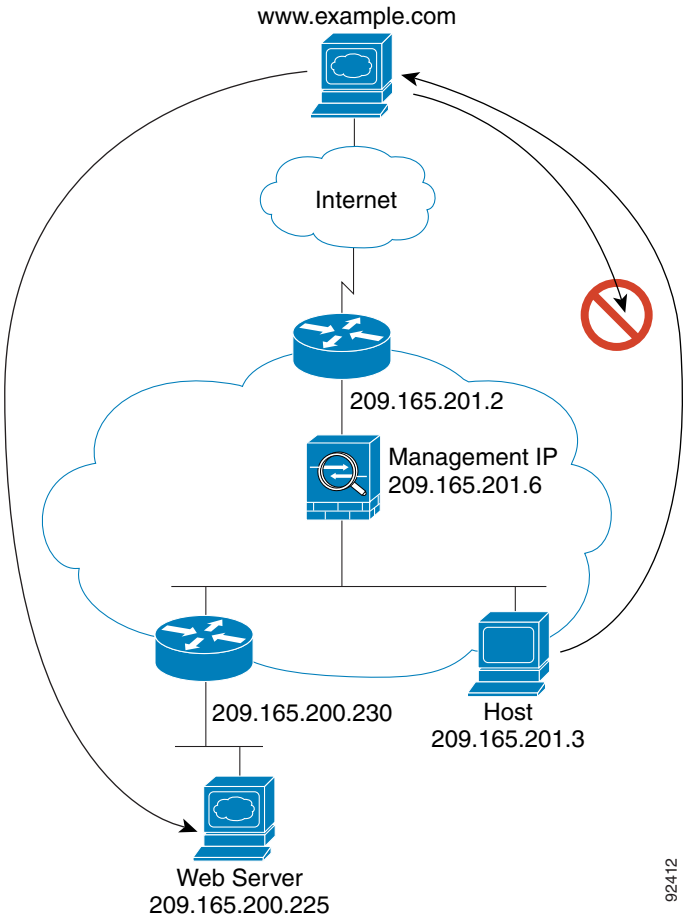

This section describes how data moves through the ASA and includes the following topics:

- An Inside User Visits a Web Server, page 6-24
- An Inside User Visits a Web Server Using NAT, page 6-25
- An Outside User Visits a Web Server on the Inside Network, page 6-26
- An Outside User Attempts to Access an Inside Host, page 6-27

An Inside User Visits a Web Server

Figure 6-9 shows an inside user accessing an outside web server.

Figure 6-9 Inside to Outside

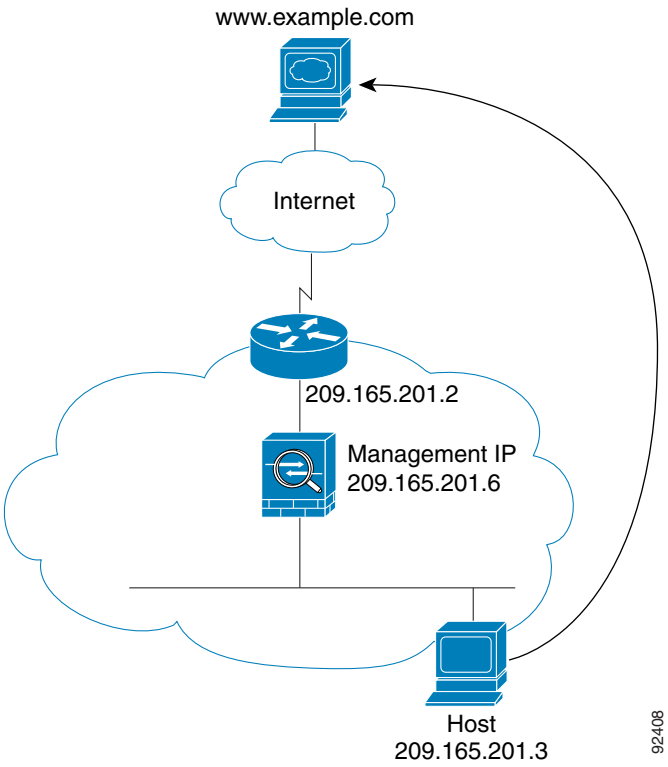

The following steps describe how data moves through the ASA (see Figure 6-9):

1. The user on the inside network requests a web page from www.example.com.
2. The ASA 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 ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA 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 ASA table, the ASA attempts to discover the MAC address by sending an ARP request or 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 ASA forwards the packet to the inside user.

An Inside User Visits a Web Server Using NAT

Figure 6-10 shows an inside user accessing an outside web server.

Figure 6-10 Inside to Outside with NAT

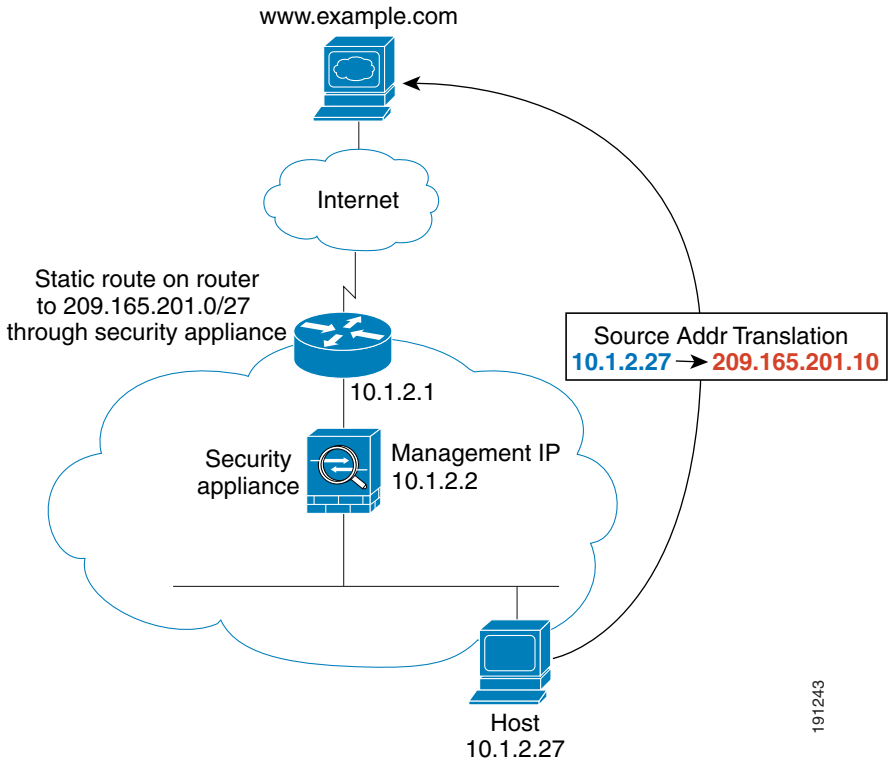

The following steps describe how data moves through the ASA (see Figure 6-10):

1. The user on the inside network requests a web page from www.example.com.
2. The ASA 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 ASA first classifies the packet according to a unique interface.
3. The ASA 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 ASA.
4. The ASA 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 ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 10.1.2.1.

   If the destination MAC address is not in the ASA table, the ASA 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 ASA performs NAT by untranslating the mapped address to the real address, 10.1.2.27.

An Outside User Visits a Web Server on the Inside Network

Figure 6-11 shows an outside user accessing the inside web server.

Figure 6-11 Outside to Inside

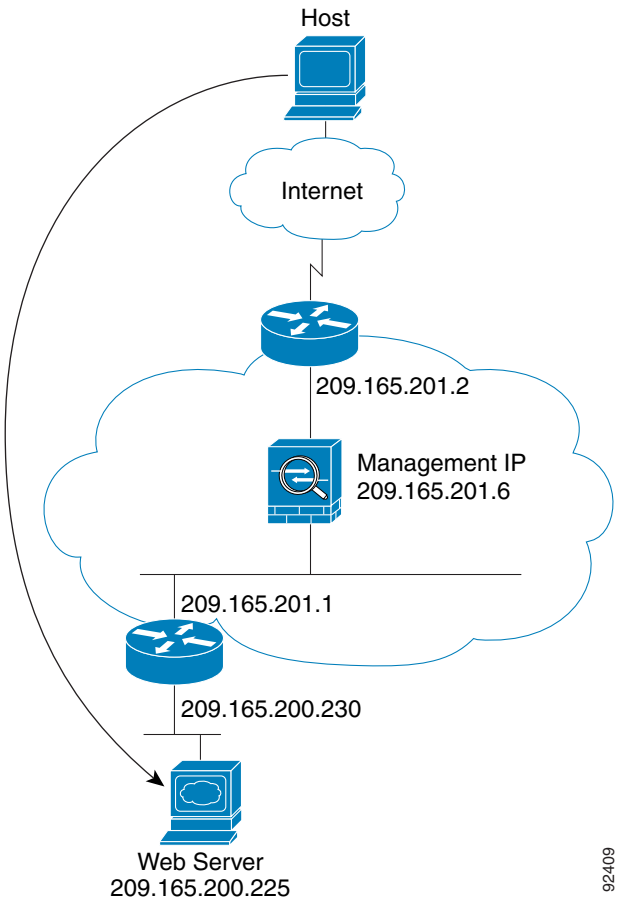

The following steps describe how data moves through the ASA (see Figure 6-11):

1. A user on the outside network requests a web page from the inside web server.
2. The ASA 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 ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.

   If the destination MAC address is not in the ASA table, the ASA 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 ASA forwards the packet to the outside user.

An Outside User Attempts to Access an Inside Host

Figure 6-12 shows an outside user attempting to access a host on the inside network.

Figure 6-12 Outside to Inside

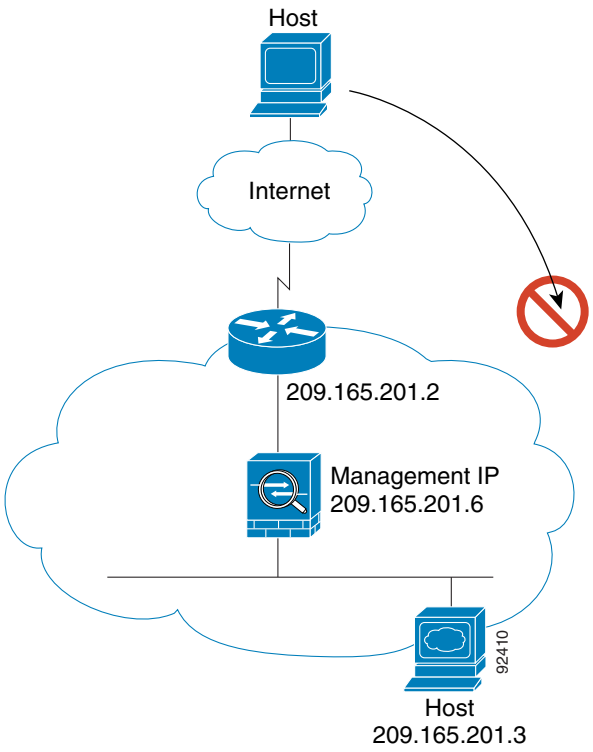

The following steps describe how data moves through the ASA (see Figure 6-12):

1. A user on the outside network attempts to reach an inside host.
2. The ASA 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 ASA first classifies the packet to a context.
3. The packet is denied because there is no access list permitting the outside host, and the ASA drops the packet.
4. If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

Using the Startup Wizard

The ASDM Startup Wizard guides you through the initial configuration of the ASA, and helps you define basic settings.

This chapter includes the following sections:

- Accessing the Startup Wizard, page 7-1
- Licensing Requirements for the Startup Wizard, page 7-1
- Guidelines and Limitations, page 7-1
- Startup Wizard Screens, page 7-2
- Feature History for the Startup Wizard, page 7-6

Accessing 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**.

Licensing Requirements for the Startup Wizard

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 mode and within a context in multiple context mode. This wizard is not supported in the System.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

IPv6 Guidelines

Supports IPv6.

Startup Wizard Screens

The actual sequence of screens is determined by your specified configuration selections. Each screen is available for all modes or models unless otherwise noted. This section includes the following topics:

- Starting Point or Welcome, page 7-2
- Basic Configuration, page 7-3
- Interface Screens, page 7-3
- Static Routes, page 7-4
- Easy VPN Remote Configuration (ASA 5505, Single Mode, Routed Mode), page 7-4
- DHCP Server, page 7-4
- Address Translation (NAT/PAT), page 7-5
- Administrative Access, page 7-5
- IPS Basic Configuration, page 7-5
- Time Zone and Clock Configuration, page 7-5
- Auto Update Server (Single Mode), page 7-5
- Startup Wizard Summary, page 7-6

Starting Point or Welcome

- To change the existing configuration, click the **Modify existing configuration** radio button.
- To set the configuration to the factory default values, click the **Reset configuration to factory defaults** radio button.
  - To configure the IP address and subnet mask of the Management 0/0 (ASA 5510 and higher) or VLAN 1 (ASA 5505) interface to be different from the default value (192.168.1.1), check the **Configure the IP address of the management interface** check box.

Note If you reset the configuration to factory defaults, you cannot undo these changes by clicking **Cancel** or by closing this screen.

In multiple context mode, this screen does not contain any parameters.

Basic Configuration

- (ASA 5505) To specify a group of configuration settings for a remote worker, check the **Configure the device for Teleworker usage** check box. See the "Easy VPN Remote Configuration (ASA 5505, Single Mode, Routed Mode)" section on page 7-4 for more information.
- For information about the hostname, domain name, and enable password, see the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1.

Interface Screens

The interface screens depend on the mode and model. This section includes the following topics:

- Interface Selection (ASA 5505), page 7-3
- Switch Port Allocation (ASA 5505), page 7-3
- Interface IP Address Configuration (ASA 5505, Routed Mode), page 7-3
- Interface Configuration - PPPoE (ASA 5505, Routed Mode, Single Mode), page 7-3
- Outside Interface Configuration - PPPoE (ASA 5510 and Higher, Routed Mode, Single Mode), page 7-4
- Management IP Address Configuration (Transparent Mode), page 7-4
- Other Interfaces Configuration (ASA 5510 and Higher), page 7-4

Interface Selection (ASA 5505)

This screen lets you 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, Inside, or DMZ (DMZ is available in routed mode only). A DMZ is a separate network located in the neutral zone between a private (inside) network and a public (outside) network.

See the "Configuring VLAN Interfaces" section on page 16-6 for more information.

Switch Port Allocation (ASA 5505)

This screen lets you allocate switch ports to Outside, Inside, or DMZ interfaces (DMZ is only available in routed mode). By default, all switch ports are assigned to VLAN 1 (Inside).

See the "Configuring VLAN Interfaces" section on page 16-6 for more information.

Interface IP Address Configuration (ASA 5505, Routed Mode)

Configure the IP address of each VLAN interface. See the "Configuring General Interface Parameters" section on page 17-7 for more information..

Interface Configuration - PPPoE (ASA 5505, Routed Mode, Single Mode)

Configure the PPoE settings for each interface. See the "PPPoE IP Address and Route Settings" section on page 17-11 for more information.

Outside Interface Configuration (ASA 5510 and Higher, Routed Mode)

- Configure the IP address of the outside interface (the interface with the lowest security level). See the "Configuring General Interface Parameters" section on page 17-7 for more information..
- To configure the IPv6 address, see the "Configuring IPv6 Addressing" section on page 17-15.

Outside Interface Configuration - PPPoE (ASA 5510 and Higher, Routed Mode, Single Mode)

Configure the PPoE settings for the outside interface. See the "PPPoE IP Address and Route Settings" section on page 17-11 for more information.

Management IP Address Configuration (Transparent Mode)

For IPv4, a management IP address is required for each bridge group for both management traffic and for traffic to pass through the ASA. This screen sets the IP address for BVI 1.

See the "Configuring Bridge Groups" section on page 18-8 for more information.

Other Interfaces Configuration (ASA 5510 and Higher)

- You can configure parameters for other interfaces. See the "Configuring General Interface Parameters" section on page 17-7 for more information.
- See the "Allowing Same Security Level Communication" section on page 17-20 for information about the Enable traffic between... check boxes.

Static Routes

Configure static routes. See Chapter 31, "Configuring Static and Default Routes," for more information.

Note For the ASA 5505, to access this screen, you must have checked the **Configure the device for Teleworker usage** check box in Basic Configuration.

Easy VPN Remote Configuration (ASA 5505, Single Mode, Routed Mode)

The ASA can act as an Easy VPN remote device to enable deployment of VPNs to remote locations. See the "Easy VPN Remote" section on page 87-105.

Note To access this screen, you must have checked the **Configure the device for Teleworker usage** check box in Basic Configuration and unchecked the **Enable Auto Update** check box in Auto Update Server (Single Mode).

DHCP Server

Configure the DHCP server. See the "Configuring a DHCP Server" section on page 23-3 for more information.

Address Translation (NAT/PAT)

Configures NAT or PAT for inside addresses (the interface with the highest security level) when accessing the outside (the interface with the lowest security level). See the "Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool" section on page 40-4 or the "Configuring Dynamic PAT (Hide)" section on page 40-8 for more information.

Administrative Access

- Configures ASDM, Telnet, or SSH access. See the "Configuring Management Access" section on page 52-3 for more information.
- To enable a secure connection to an HTTP server to access ASDM, check the **Enable HTTP server for HTTPS/ASDM access** check box. See the "Configuring Management Access" section on page 52-3 for more information.
- To allow ASDM to collect and display statistics, check the **Enable ASDM history metrics** check box. See the "Enabling History Metrics" section on page 4-34 for more information.

IPS Basic Configuration

In single context mode, you can use the Startup Wizard in ASDM to configure basic IPS network configuration. These settings are saved to the IPS configuration, not the ASA configuration. See the "Configuring Basic IPS Module Network Settings" section on page 81-12.

Time Zone and Clock Configuration

Configure the clock parameters. See the "Setting the Date and Time" section on page 21-2 for more information.

Auto Update Server (Single Mode)

- Configure an auto update server by checking the **Enable Auto Update Server for ASA** check box. See the "Configuring Auto Update" section on page 101-30 for more information.
- If you have an IPS module, you can check the **Enable Signature and Engine Updates from Cisco.com** check box. Set the following additional parameters:
  - Enter your Cisco.com username and password, and then confirm the password.
  - Enter the start time in hh:mm:ss format, using a 24-hour clock.

Note For the ASA 5505, to access this screen, you must have checked the **Configure the device for Teleworker usage** check box in Basic Configuration.

Startup Wizard Summary

This screen summarizes all of the configuration settings that you have made for the ASA.

- To change any of the settings in previous screens, click **Back**.
- 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 ASA 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**.

Feature History for the Startup Wizard

Table 7-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 7-1 Feature History for the Startup Wizard

Feature Name	Platform Releases	Feature Information
Startup Wizard	7.0(1)	This feature was introduced. We introduced the Wizards > Startup Wizard screen.
IPS Configuration	8.4(1)	For the IPS module, the IPS Basic Configuration screen was added to the startup wizard. Signature updates for the IPS module were also added to the Auto Update screen. The Time Zone and Clock Configuration screen was added to ensure the clock is set on the ASA; the IPS module gets its clock from the ASA. We introduced or modified the following screens: Wizards > Startup Wizard > IPS Basic Configuration Wizards > Startup Wizard > Auto Update Wizards > Startup Wizard > Time Zone and Clock Configuration

PART 1

Configuring High Availability and Scalability

Configuring Multiple Context Mode

This chapter describes how to configure multiple security contexts on the ASA and includes the following sections:

- Information About Security Contexts, page 9-1
- Licensing Requirements for Multiple Context Mode, page 9-13
- Guidelines and Limitations, page 9-14
- Default Settings, page 9-15
- Configuring Multiple Contexts, page 9-15
- Changing Between Contexts and the System Execution Space, page 9-25
- Managing Security Contexts, page 9-26
- Monitoring Security Contexts, page 9-30
- Feature History for Multiple Context Mode, page 9-33

Information About Security Contexts

You can partition a single ASA into multiple virtual devices, known as security contexts. Each context acts as an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. For unsupported features in multiple context mode, see the "Guidelines and Limitations" section on page 9-14.

This section provides an overview of security contexts and includes the following topics:

- Common Uses for Security Contexts, page 9-2
- Context Configuration Files, page 9-2
- How the ASA Classifies Packets, page 9-3
- Cascading Security Contexts, page 9-6
- Management Access to Security Contexts, page 9-7
- Information About Resource Management, page 9-8
- Information About MAC Addresses, page 9-11

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 ASA, 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 ASA.

Context Configuration Files

This section describes how the ASA implements multiple context mode configurations and includes the following topics:

- Context Configurations, page 9-2
- System Configuration, page 9-2
- Admin Context Configuration, page 9-2

Context Configurations

For each context, the ASA includes a configuration that identifies the security policy, interfaces, and all the options you can configure on a standalone device. You can store context configurations in flash memory, or you can download them from a TFTP, FTP, or HTTP(S) server.

System Configuration

The system administrator adds and manages contexts by configuring each context configuration location, allocated interfaces, and other context operating parameters in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the ASA. 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 system configuration does include a specialized failover interface for failover traffic only.

Admin Context Configuration

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. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users. The admin context must reside on flash memory, and not remotely.

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 ASA Classifies Packets

Each packet that enters the ASA must be classified, so that the ASA can determine to which context to send a packet. This section includes the following topics:

- Valid Classifier Criteria, page 9-3
- Classification Examples, page 9-4

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 9-3
- Unique MAC Addresses, page 9-3
- NAT Configuration, page 9-3

Note For management traffic destined for an interface, the interface IP address is used for classification.

The routing table is not used for packet classification.

Unique Interfaces

If only one context is associated with the ingress interface, the ASA 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 unique MAC addresses assigned to the interface in each context. An upstream router cannot route directly to a context without unique MAC addresses. By default, auto-generation of MAC addresses is enabled. You can also set the MAC addresses manually when you configure each interface.

NAT Configuration

If you disable use of unique MAC addresses, then the ASA uses the mapped addresses in your NAT configuration to classify packets. We recommend using MAC addresses instead of NAT, so that traffic classification can occur regardless of the completeness of the NAT configuration.

Classification Examples

Figure 9-1 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 9-1 Packet Classification with a Shared Interface Using MAC Addresses

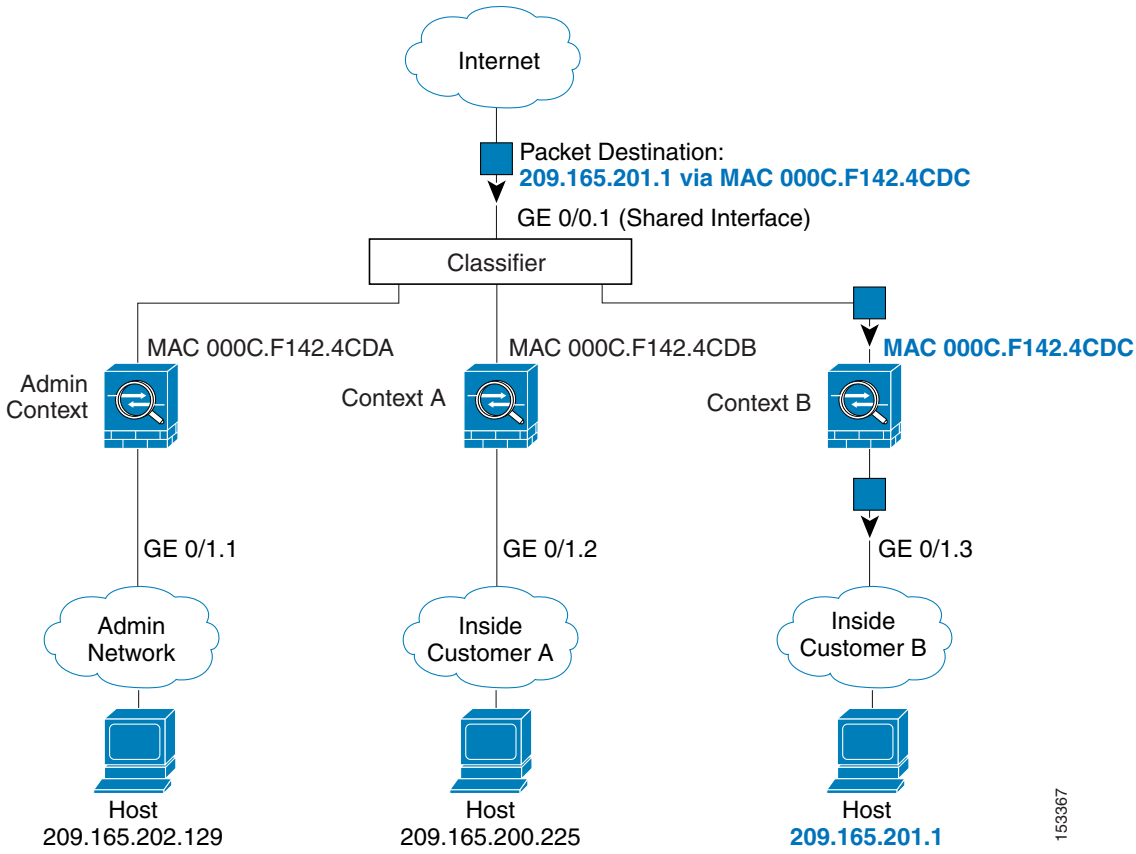

Note that all new incoming traffic must be classified, even from inside networks. Figure 9-2 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.

Figure 9-2 Incoming Traffic from Inside Networks

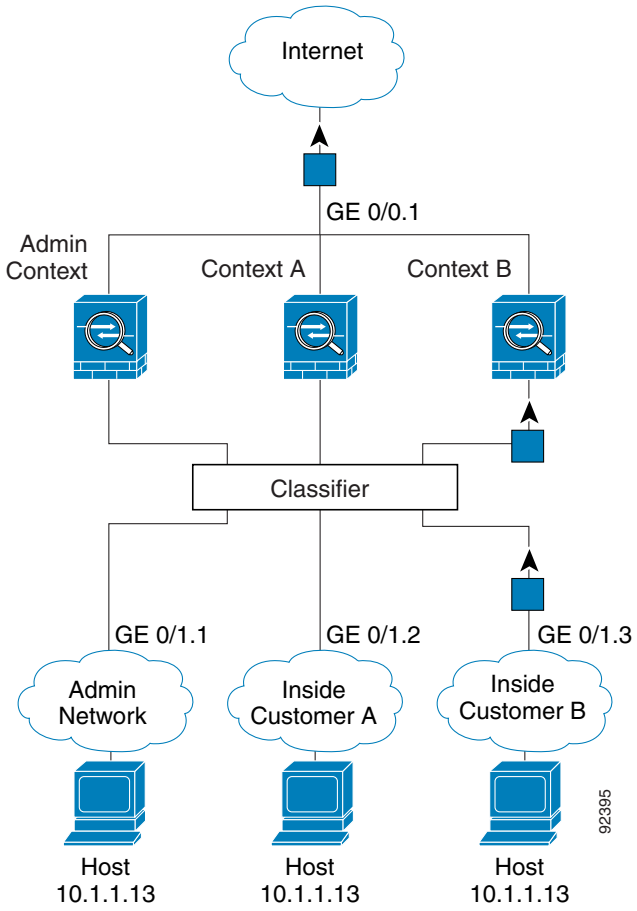

For transparent firewalls, you must use unique interfaces. Figure 9-3 shows a packet destined to a host on the Context B inside network from the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

Figure 9-3 Transparent Firewall Contexts

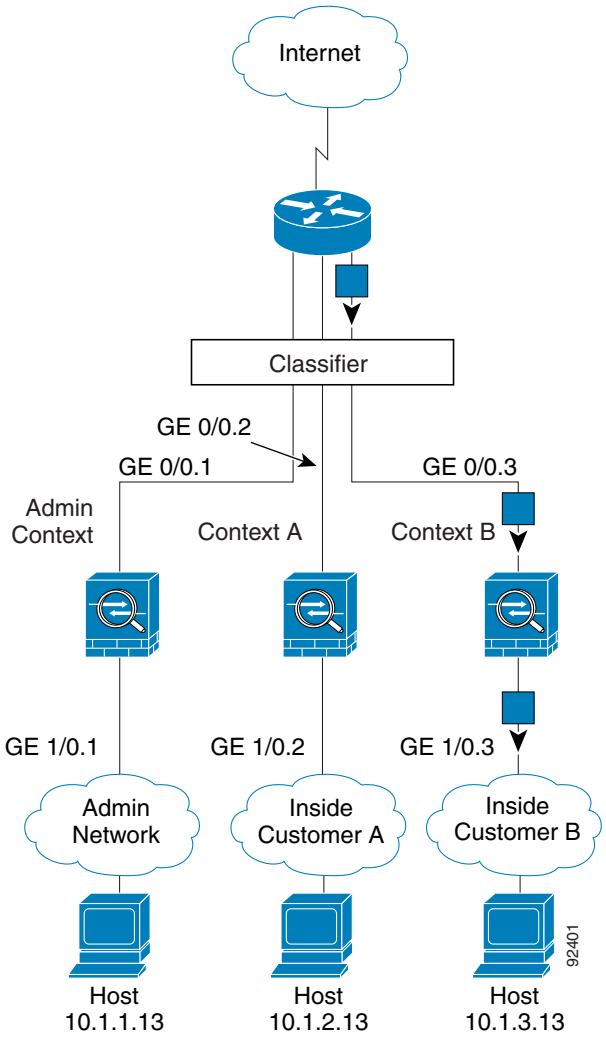

Cascading Security Contexts

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 unique MAC addresses for each context interface (the default setting). 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 9-4 shows a gateway context with two contexts behind the gateway.

Figure 9-4 Cascading Contexts

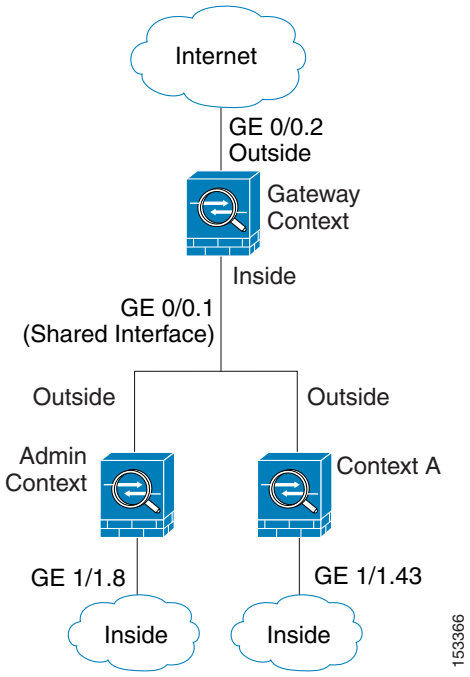

Management Access to Security Contexts

The ASA 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 context administrator:

- System Administrator Access, page 9-7
- Context Administrator Access, page 9-8

System Administrator Access

You can access the ASA as a system administrator in two ways:

- Access the ASA console.

  From the console, you access the *system execution space*, which means that any commands you enter affect only the system configuration or the running of the system (for run-time commands).

- Access the admin context using Telnet, SSH, or ASDM.

  See Chapter 52, "Configuring Management Access," 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. To log in with a new 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 to enable_15, so you must log in again as "admin" by entering the **login** command. When you change to context B, you must again enter the **login** command to log in as "admin."

The system execution space does not support any AAA commands, but you can configure its own enable password, as well as usernames in the local database to provide individual logins.

Context Administrator Access

You can access a context using Telnet, SSH, or ASDM. If you log in to a non-admin context, you can only access the configuration for that context. You can provide individual logins to the context. See Chapter 52, "Configuring Management Access," to enable Telnet, SSH, and ASDM access and to configure management authentication.

Information About Resource Management

By default, all security contexts have unlimited access to the resources of the ASA, except where maximum limits per context are enforced; the only exception is VPN resources, which are disabled by default. 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. For VPN resources, you must configure resource management to allow any VPN tunnels.

This section includes the following topics:

- Resource Classes, page 9-8
- Resource Limits, page 9-8
- Default Class, page 9-9
- Using Oversubscribed Resources, page 9-10
- Using Unlimited Resources, page 9-11

Resource Classes

The ASA manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. 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.

Resource Limits

You can set the limit for individual resources as a percentage (if there is a hard system limit) or as an absolute value.

For most resources, the ASA does not set aside a portion of the resources for each context assigned to the class; rather, the ASA 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. The exception is VPN resource types, which you cannot oversubscribe, so the resources assigned to each context are guaranteed. To accommodate temporary bursts of VPN sessions beyond the amount assigned, the ASA supports a "burst" VPN resource type, which is equal to the remaining unassigned VPN sessions. The burst sessions *can* be oversubscribed, and are available to contexts on a first-come, first-served basis.

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.

For most resources, the default class provides unlimited access to resources for all contexts, except for the following limits:

- Telnet sessions—5 sessions. (The maximum per context.)
- SSH sessions—5 sessions. (The maximum per context.)
- IPsec sessions—5 sessions. (The maximum per context.)
- MAC addresses—65,535 entries. (The maximum per context.)
- VPN site-to-site tunnels—0 sessions. (You must manually configure the class to allow any VPN sessions.)

Figure 9-5 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 9-5 Resource Classes

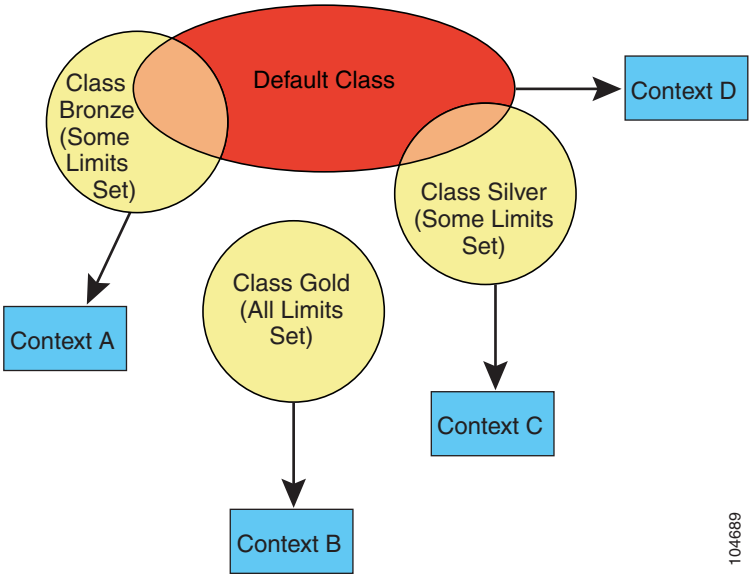

Using Oversubscribed Resources

You can oversubscribe the ASA by assigning more than 100 percent of a resource across all contexts (with the exception of non-burst VPN resources). 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 9-6.)

Figure 9-6 Resource Oversubscription

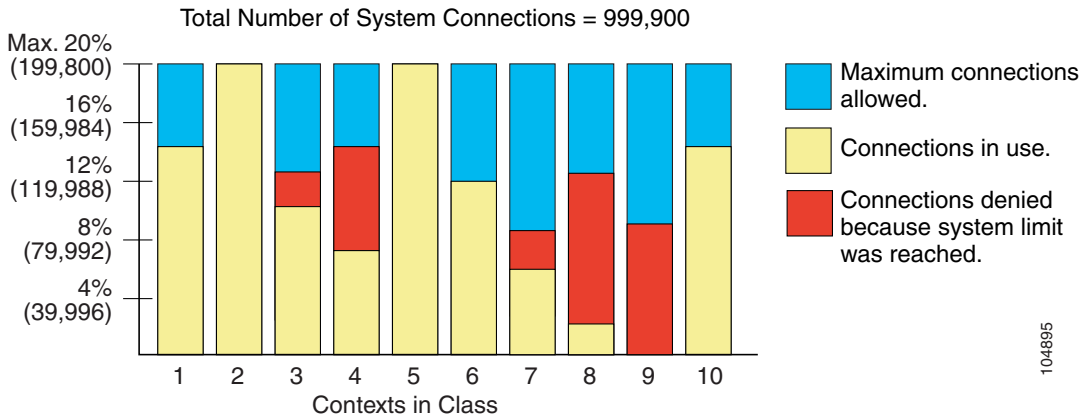

Using Unlimited Resources

The ASA 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. 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 9-7.) Setting unlimited access is similar to oversubscribing the ASA, except that you have less control over how much you oversubscribe the system.

Figure 9-7 Unlimited Resources

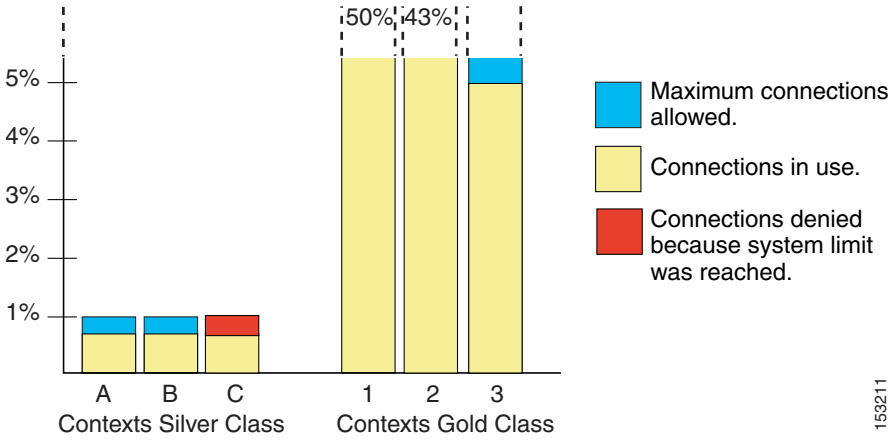

Information About MAC Addresses

To allow contexts to share interfaces, you should 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 other classification methods are attempted that might not provide full coverage. See the "How the ASA Classifies Packets" section on page 9-3 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 "Configuring the MAC Address and MTU" section on page 17-13 to manually set the MAC address.

This section includes the following topics:

- Default MAC Address, page 9-12
- Interaction with Manual MAC Addresses, page 9-12
- Failover MAC Addresses, page 9-12
- MAC Address Format, page 9-12

Default MAC Address

If you disable MAC address generation, see the following default MAC addresses:

- For the ASA 5500 series appliances—The physical interface uses the burned-in MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.
- For the ASASM—All VLAN interfaces use the same MAC address, derived from the backplane MAC address.

See the following sections for your release for additional information about automatic MAC address generation. See also the "MAC Address Format" section on page 9-12.

8.6(1) and Later

Automatic MAC address generation is enabled—Uses an autogenerated prefix. The ASA autogenerates the prefix based on the last two bytes of the interface (ASA 5500) or backplane (ASASM) MAC address. You cannot use the legacy auto-generation method (without a prefix).

Note To maintain hitless upgrade for failover pairs, the ASA does not convert an existing auto-generation configuration upon a reload if failover is enabled. However, we strongly recommend that you manually change to the prefix method of generation when using failover, especially for the ASASM. Without the prefix method, ASASMs installed in different slot numbers experience a MAC address change upon failover, and can experience traffic interruption. After upgrading, to use the prefix method of MAC address generation, reenable MAC address autogeneration to use a prefix.

8.5(1) (ASASM only)

Automatic MAC address generation is enabled—Uses the legacy auto-generation method (without a prefix).

Earlier Releases

Automatic MAC address generation is disabled.

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 (when using a prefix) 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 ASA 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 on page 9-12 section for more information.

MAC Address Format

The MAC address format without a prefix is a legacy version not supported on newer ASA versions.

MAC Address Format Using a Prefix

The ASA generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or an autogenerated prefix based on the last two bytes of the interface (ASA 5500) or backplane (ASASM) MAC address, and *zz.zzzz* is an internal counter generated by the ASA. 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 ASA converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the ASA native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

MAC Address Format Without a Prefix (Legacy Method; Not Available in 8.6(1) and Later)

Without a prefix, the MAC address is generated using the following format:

- Active unit MAC address: 12*_slot.port_subid.contextid*.
- Standby unit MAC address: 02*_slot.port_subid.contextid*.

For platforms with no interface slots, the slot is always 0. The *port* is the interface port. The *subid* is an internal ID for the subinterface, which is not viewable. The *contextid* is an internal ID for the context, viewable with the **show context detail** command. For example, the interface GigabitEthernet 0/1.200 in the context with the ID 1 has the following generated MAC addresses, where the internal ID for subinterface 200 is 31:

- Active: 1200.0131.0001
- Standby: 0200.0131.0001

This MAC address generation method does not allow for persistent MAC addresses across reloads, does not allow for multiple ASAs on the same network segment (because unique MAC addresses are not guaranteed), and does not prevent overlapping MAC addresses with manually assigned MAC addresses. We recommend using a prefix with the MAC address generation to avoid these issues.

Licensing Requirements for Multiple Context Mode

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.*

Model	License Requirement
ASA 5550	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, or 100 contexts.*
ASA 5580	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.*
ASA 5512-X	No support.
ASA 5515-X	Security Plus License: 2 contexts. *Optional license: 5 contexts.*
ASA 5525-X	Base License: 2 contexts. *Optional licenses: 5, 10, or 20 contexts.*
ASA 5545-X	Base License: 2 contexts. *Optional licenses: 5, 10, 20, or 50 contexts.*
ASA 5555-X	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, or 100 contexts.*
ASA 5585-X with SSP-10	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, or 100 contexts.*
ASA 5585-X with SSP-20, -40, and -60	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.*
ASASM	Base License: 2 contexts. *Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.*

Prerequisites

After you are in multiple context mode, connect tothe admin context to access the system configuration. You cannot configure the system from a non-admin context. By default, after you enable multiple context mode, you can connect to the admin context by using the default management IP address. See Chapter 3, "Getting Started," for more information about connecting to the ASA.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode; set the firewall mode per context.

Failover Guidelines

Active/Active mode failover is only supported in multiple context mode.

IPv6 Guidelines

Supports IPv6.

Model Guidelines

Does not support the ASA 5505 and 5512-X.

Unsupported Features

Multiple context mode does not support the following features:

- RIP
- OSPFv3. (OSPFv2 is supported.)
- Multicast routing
- Threat Detection
- Unified Communications
- QoS
- Remote access VPN. (Site-to-site VPN is supported.)

Additional Guidelines

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.

Default Settings

- By default, the ASA is in single context mode.
- See the "Default Class" section on page 9-9.
- See the "Default MAC Address" section on page 9-12.

Configuring Multiple Contexts

This section describes how to configure multiple context mode and includes the following topics:

- Task Flow for Configuring Multiple Context Mode, page 9-15
- Enabling or Disabling Multiple Context Mode, page 9-16
- Configuring a Class for Resource Management, page 9-17
- Configuring a Security Context, page 9-20
- Automatically Assigning MAC Addresses to Context Interfaces, page 9-24

Task Flow for Configuring Multiple Context Mode

To configure multiple context mode, perform the following steps:

Step 1 Enable multiple context mode. See the "Enabling or Disabling Multiple Context Mode" section on page 9-16.

Step 2 (Optional) Configure classes for resource management. See the "Configuring a Class for Resource Management" section on page 9-17.

Step 3 Configure interfaces in the system execution space.

- ASA 5500—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
- ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

Step 4 Configure security contexts. See the "Configuring a Security Context" section on page 9-20.

Step 5 (Optional) Customize MAC address assignments. See the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 9-24.

Step 6 Complete interface configuration in the context. See Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Enabling or Disabling Multiple Context Mode

Your ASA might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you need to convert from single mode to multiple mode, follow 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 Chapter 13, "Configuring Active/Active Failover," 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 using the CLI; because changing modes requires confirmation, you cannot use the Command Line Interface tool. This section describes changing modes at the CLI.

This section includes the following topics:

- Enabling Multiple Context Mode, page 9-16
- Restoring Single Context Mode, page 9-17

Enabling Multiple Context Mode

When you convert from single mode to multiple mode, the ASA 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 ASA automatically adds an entry for the admin context to the system configuration with the name "admin."

Prerequisites

Back up your startup configuration. When you convert from single mode to multiple mode, the ASA converts the running configuration into two files. The original startup configuration is not saved. See the "Managing Files" section on page 101-14.

Detailed Steps

Command	Purpose
mode multiple **Example:** `hostname(config)# mode multiple`	Changes to multiple context mode. You are prompted to reboot the ASA.

Restoring Single Context Mode

To copy the old running configuration to the startup configuration and to change the mode to single mode, perform the following steps.

Prerequisites

Perform this procedure in the system execution space.

Detailed Steps

	Command	Purpose
Step 1	**copy disk0:old_running.cfg startup-config** **Example:** `hostname(config)# copy disk0:old_running.cfg startup-config`	Copies the backup version of your original running configuration to the current startup configuration.
Step 2	**mode single** **Example:** `hostname(config)# mode single`	Sets the mode to single mode. You are prompted to reboot the ASA.

Configuring a Class for Resource Management

To configure a class in the system configuration, perform the following steps. You can change the value of a particular resource limit by reentering the command with a new value.

Prerequisites

Perform this procedure in the system execution space.

Guidelines

Table 9-1 lists the resource types and the limits.

Table 9-1 Resource Names and Limits

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit[1]	Description
ASDM Sessions	Concurrent	1 minimum 5 maximum	32	ASDM management sessions. **Note** 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.
Connections Conns/sec	Concurrent or Rate	N/A	Concurrent connections: See the "Supported Feature Licenses Per Model" section on page 5-1 for the connection limit available for your model. Rate: N/A	TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts.
Hosts	Concurrent	N/A	N/A	Hosts that can connect through the ASA.
Inspects/sec	Rate	N/A	N/A	Application inspections per second.
MAC Entries	Concurrent	N/A	65,535	For transparent firewall mode, the number of MAC addresses allowed in the MAC address table.
Routes	Concurrent	N/A	N/A	Dynamic routes.
Site-to-Site VPN Burst	Concurrent	N/A	The Other VPN session amount for your model minus the sum of the sessions assigned to all contexts for Site-to-Site VPN.	The number of site-to-site VPN sessions allowed beyond the amount assigned to a context with Site-to-Site VPN. For example, if your model supports 5000 sessions, and you assign 4000 sessions across all contexts with Site-to-Site VPN, then the remaining 1000 sessions are available for Site-to-Site VPN Burst. Unlike Site-to-Site VPN, which guarantees the sessions to the context, Site-to-Site VPN Burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
Site-to-Site VPN	Concurrent	N/A	See the "Supported Feature Licenses Per Model" section on page 5-1 for the Other VPN sessions available for your model.	Site-to-site VPN sessions. You cannot oversubscribe this resource; all context assignments combined cannot exceed the model limit. The sessions you assign for this resource are guaranteed to the context.
SSH	Concurrent	1 minimum 5 maximum	100	SSH sessions.

Table 9-1 Resource Names and Limits (continued)

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit[1]	Description
Syslogs/sec	Rate	N/A	N/A	Syslog messages per second.
Telnet	Concurrent	1 minimum 5 maximum	100	Telnet sessions.
xlates	Concurrent	N/A	N/A	Network address translations.

1. If this column value is N/A, then you cannot set a percentage of the resource because there is no hard system limit for the resource.

Detailed 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 **Configuration > Context Management > Resource Class**, and click **Add**.

The Add Resource Class dialog box appears.

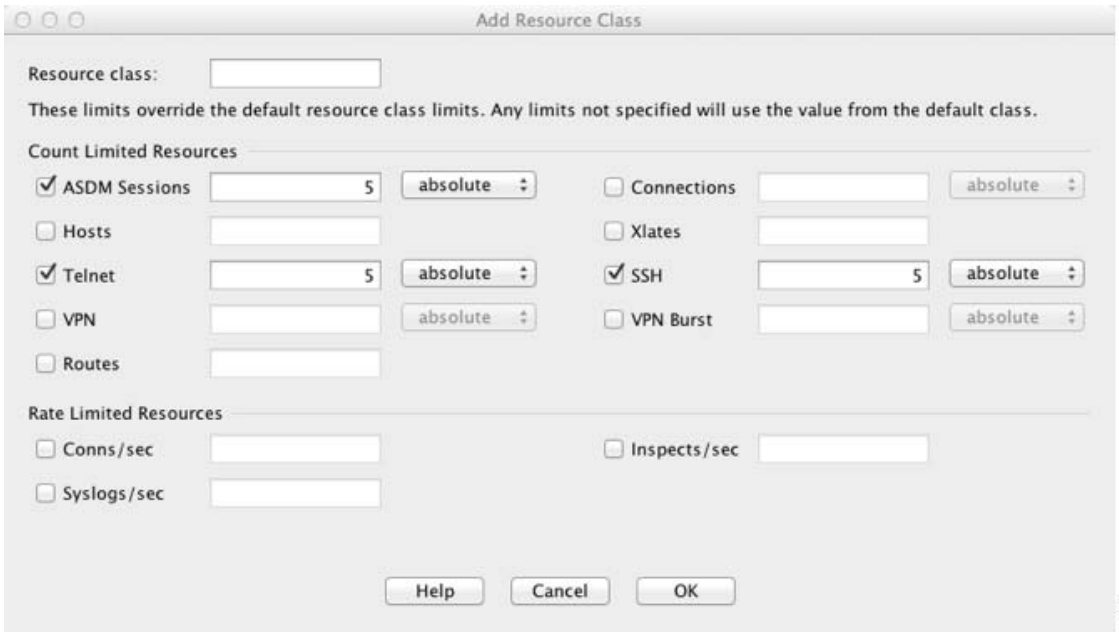

Step 3 In the Resource Class field, enter a class name up to 20 characters in length.

Step 4 In the Count Limited Resources area, set the concurrent limits for resources.

See the Table 9-1 on page 9-18 for a description of each resource type.

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. For most resources, 0 sets the limit to unlimited. For VPN types, 0 sets the limit none.

Step 5 In the Rate Limited Resources area, set the rate limit for resources.

See the Table 9-1 on page 9-18 for a description of each resource type.

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. 0 sets the limit to unlimited

Step 6 Click **OK**.

Configuring a Security Context

The security context definition in the system configuration identifies the context name, configuration file URL, interfaces that a context can use, and other settings.

Prerequisites

- Perform this procedure in the system execution space.
- For the ASASM, assign VLANs to the ASASM on the switch according to Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- For the ASA 5500, configure physical interface parameters, VLAN subinterfaces, EtherChannels, and redundant interfaces according to Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."

Detailed 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 **Configuration > Context Management > Security Contexts**, and click **Add**.

The Add Context dialog box appears.

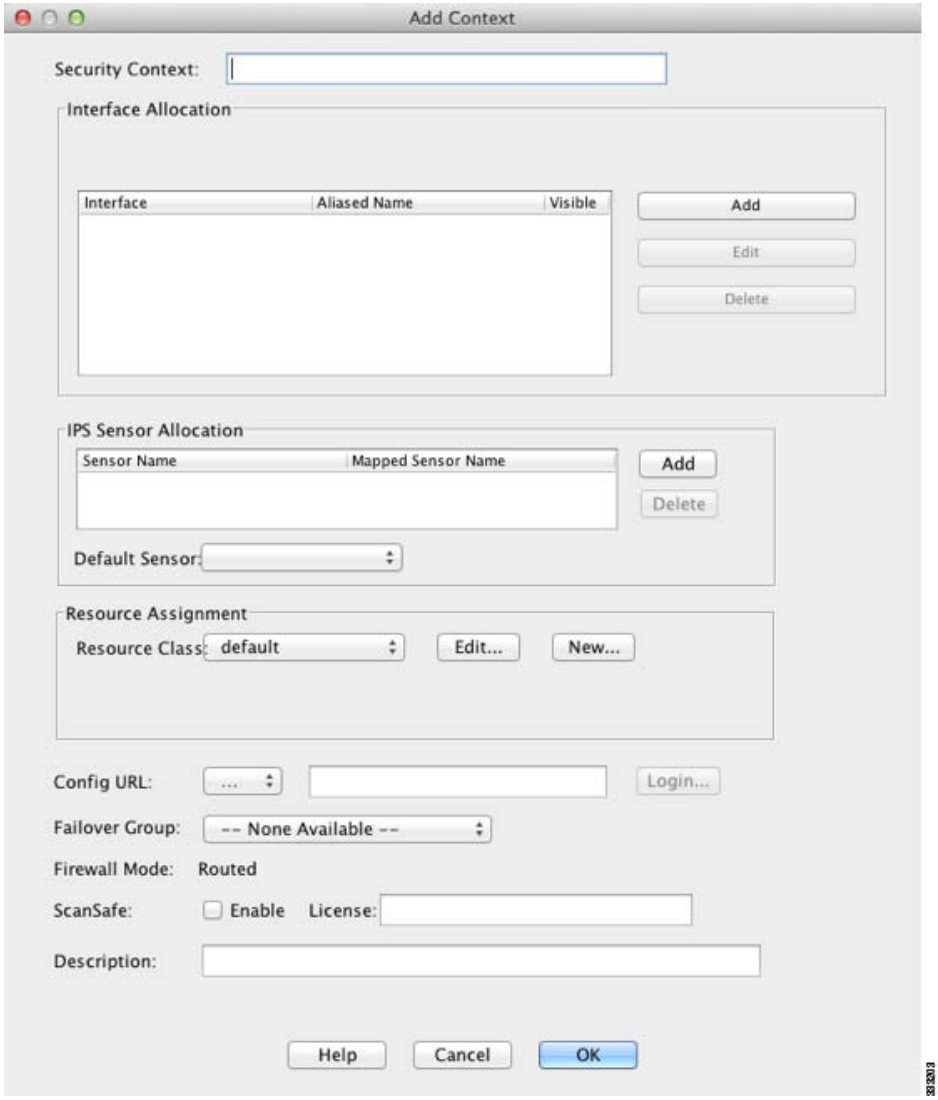

Step 3 In the Security Context field, enter the context name as a string up to 32 characters long.

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.

a. 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.

b. (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.

a. (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.

   – 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.

   – (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.

b. (Optional) To enable context users to see physical interface properties even if you set an aliased name, check **Show Hardware Properties in Context**.

c. Click **OK** to return to the Add Context dialog box.

Step 5 (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 81, “Configuring the ASA IPS Module.”

Step 6 (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 a Class for Resource Management" section on page 9-17 for more information.

Step 7 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

a. (Optional) For external file systems, set the username and password by clicking **Login**.

Step 8 (Optional) To set the failover group for Active/Active failover, choose the group name in the Failover Group drop-down list.

Step 9 (Optional) To enable ScanSafe inspection in this context, click **Enable**. To override the license set in the system configuration, enter a license in the License field.

Step 10 (Optional) Add a description in the Description field.

Step 11 Click **OK** to return to the Security Contexts pane.

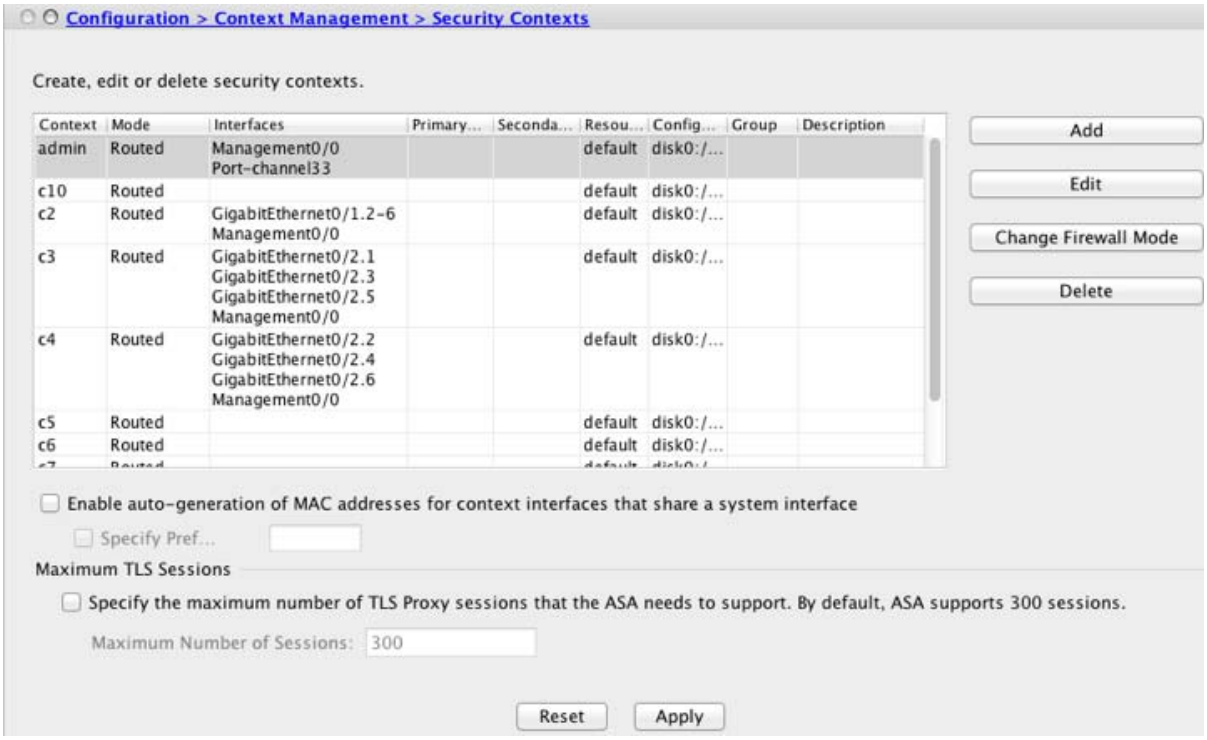

Step 12 (Optional) To set the firewall mode to transparent, select the context and click **Change Firewall Mode**.

You see the following confirmation dialog box:

If this is a new context, there is no configuration to erase. Click **Change Mode** to change to transparent firewall mode.

If this is an existing context, then be sure to back up the configuration before you change the mode.

Step 13 To customize auto-generation of MAC addresses, see the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 9-24.

Step 14 To specify the maximum TLS Proxy sessions for the device, check the **Specify the maximum number of TLS Proxy sessions that the ASA needs to support** check box. For more information about TLS proxy, see Chapter 66, "Configuring the TLS Proxy for Encrypted Voice Inspection."

Automatically Assigning MAC Addresses to Context Interfaces

This section describes how to configure auto-generation of MAC addresses.

The MAC address is used to classify packets within a context. See the "Information About MAC Addresses" section on page 9-11 for more information, especially if you are upgrading from an earlier ASA version. See also the "Viewing Assigned MAC Addresses" section on page 9-32.

Guidelines

- When you configure a name for the interface in a context, the new MAC address is generated immediately. If you enable this feature after you configure context interfaces, then MAC addresses are generated for all interfaces immediately after you enable it. If you disable this feature, 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.
- In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context. See the "Configuring the MAC Address and MTU" section on page 17-13 to manually set the MAC address.

Detailed 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 **Configuration > Context Management > Security Contexts**, and check **Mac-Address auto**. If you do not enter a prefix, then the ASA autogenerates the prefix based on the last two bytes of the interface (ASA 5500) or backplane (ASASM) MAC address.

Step 3 (Optional) Check the **Prefix** check box, and in the field, enter a decimal value between 0 and 65535.

This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address. See the "MAC Address Format" section on page 9-12 section for more information about how the prefix is used.

Changing Between Contexts and the System Execution Space

If you log in to the system execution space (or the admin context), you can change between contexts and perform configuration and monitoring tasks within each context. The running configuration that you edit in a configuration mode depends on your location. When you are in the system execution space, the running configuration consists only of the system configuration; when you are in a context, the running configuration consists only of that context.

Detailed Steps

Step 1 To configure the System, in the Device List pane, double-click **System** under the active device IP address.

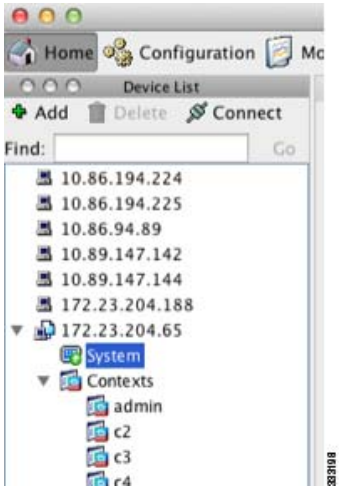

Step 2 To configure a context, in the Device List pane, double-click the context name under the active device IP address.

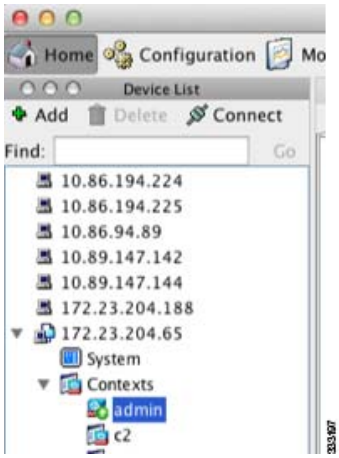

Managing Security Contexts

This section describes how to manage security contexts and includes the following topics:

- Removing a Security Context, page 9-26
- Changing the Admin Context, page 9-27
- Changing the Security Context URL, page 9-28
- Reloading a Security Context, page 9-29

Removing a Security Context

You cannot remove the current admin context.

Note If you use failover, there is a delay between when you remove the context on the active unit and when the context is removed on the standby unit.

Prerequisites

Perform this procedure in the system execution space.

Detailed 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 **Configuration > Context Management > Security Contexts**.

Step 3 Select the context you want to delete, and click **Delete**.

The Delete Context dialog box appears.

Step 4 If you might want to re-add this context later, and want to keep the configuration file for future use, uncheck the **Also delete config URL file from the disk** check box.

If you want to delete the configuration file, then leave the check box checked.

Step 5 Click **Yes**.

Changing the Admin Context

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. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users.

Note For ASDM, you cannot change the admin context within ASDM because your ASDM session would disconnect. You can perform this procedure using the Command Line Interface tool noting that you will have to reconnect to the new admin context.

Guidelines

You can set any context to be the admin context, as long as the configuration file is stored in the internal flash memory.

Prerequisites

Perform this procedure in the system execution space.

Detailed 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 **Tools > Command Line Interface**.

The Command Line Interface dialog box appears.

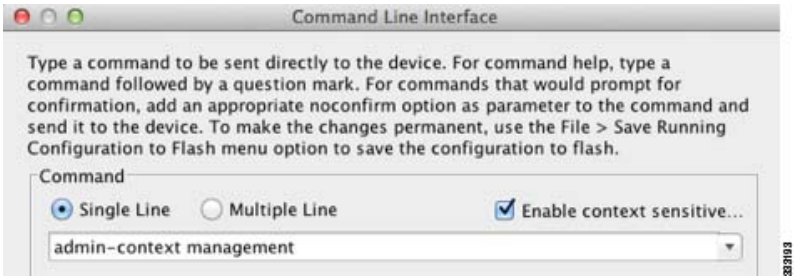

Step 3 Enter the following command:

admin-context *context_name*

Step 4 Click **Send**.

Any remote management sessions, such as Telnet, SSH, or HTTPS (ASDM), that are connected to the admin context are terminated. You must reconnect to the new admin context.

Note A few system configuration commands, including **ntp server**, identify an interface name that belongs to the admin context. If you change the admin context, and that interface name does not exist in the new admin context, be sure to update any system commands that refer to the interface.

Changing the Security Context URL

This section describes how to change the context URL.

Guidelines

- You cannot change the security context URL without reloading the configuration from the new URL. The ASA merges the new configuration with the current running configuration.
- Reentering the same URL also merges the saved configuration with the running configuration.
- A merge adds any new commands from the new configuration to the running configuration.
  - If the configurations are the same, no changes occur.
  - If commands conflict or if commands affect the running of the context, then the effect of the merge depends on the command. You might get errors, or you might have unexpected results. If the running configuration is blank (for example, if the server was unavailable and the configuration was never downloaded), then the new configuration is used.
- If you do not want to merge the configurations, you can clear the running configuration, which disrupts any communications through the context, and then reload the configuration from the new URL.

Prerequisites

Perform this procedure in the system execution space.

Detailed 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 **Configuration > Context Management > Security Contexts**.

Step 3 Select the context you want to edit, and click **Edit**.

The Edit Context dialog box appears.

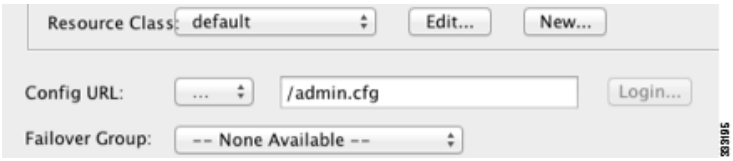

Step 4 Enter a new URL in the Config URL field, and click **OK**.

The system immediately loads the context so that it is running.

Reloading a Security Context

You can reload the context in two ways:

- Clear the running configuration and then import the startup configuration.

  This action clears most attributes associated with the context, such as connections and NAT tables.

- Remove the context from the system configuration.

  This action clears additional attributes, such as memory allocation, which might be useful for troubleshooting. However, to add the context back to the system requires you to respecify the URL and interfaces.

This section includes the following topics:

- Reloading by Clearing the Configuration, page 9-29
- Reloading by Removing and Re-adding the Context, page 9-30

Reloading by Clearing the Configuration

To reload the context by clearing the context configuration and reloading the configuration from the URL, perform the following steps.

Detailed Steps

Step 1 In the Device List pane, double-click the context name under the active device IP address.

Step 2 Choose **Tools > Command Line Interface**.

The Command Line Interface dialog box appears.

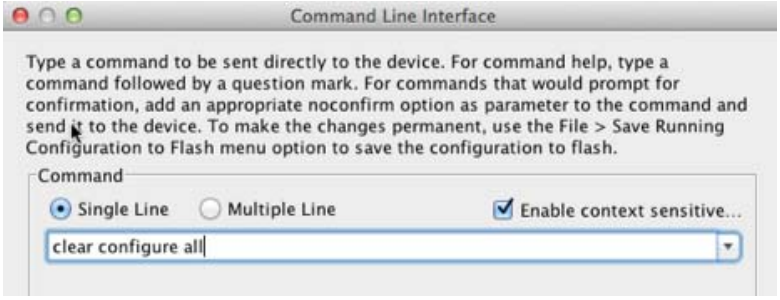

Step 3 Enter the following command:

```
clear configure all
```

Step 4 Click **Send**.

The context configuration is cleared.

Step 5 Choose **Tools > Command Line Interface** again.

The Command Line Interface dialog box appears.

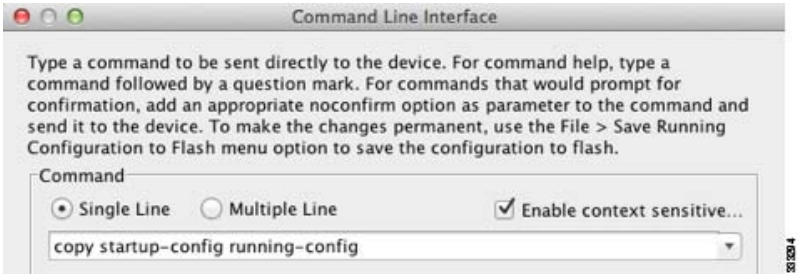

Step 6 Enter the following command:

```
copy startup-config running-config
```

Step 7 Click **Send**.

The ASA reloads the configuration. The ASA copies the configuration from the URL specified in the system configuration. You cannot change the URL from within a context.

Reloading by Removing and Re-adding the Context

To reload the context by removing the context and then re-adding it, perform the steps in the following sections:

1. "Removing a Security Context" section on page 9-26. Be sure to uncheck the **Also delete config URL file from the disk** check box.
2. "Configuring a Security Context" section on page 9-20

Monitoring Security Contexts

This section describes how to view and monitor context information and includes the following topics:

- Monitoring Context Resource Usage, page 9-31
- Viewing Assigned MAC Addresses, page 9-32

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/Telnet/SSH**—Shows the usage of ASDM, Telnet, and SSH connections.
  - Context—Shows the name of each context.

  For each access method, see the following usage statistics:
  - 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.
- **Routes**—Shows the usage of dynamic routes.
  - 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 network 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.
- VPN—Shows the usage of VPN site-to-site tunnels.
  - Context—Shows the name of each context.
  - VPN Connections—Shows usage of guaranteed VPN sessions.
  - VPN Burst Connections—Shows usage of burst VPN sessions.

    Existing (#)—Shows the number of existing tunnels.

    Peak (#)—Shows the peak number of tunnels 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.

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 9-32
- Viewing MAC Addresses Within a Context, page 9-33

Viewing MAC Addresses in the System Configuration

This section describes how to view MAC addresses in the system configuration.

Guidelines

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.

Detailed 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 **Configuration > Context Management > Security Contexts**, and view the Primary MAC and Secondary MAC columns.

Viewing MAC Addresses Within a Context

This section describes how to view MAC addresses within a context.

Detailed 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 **Configuration > Interfaces**, 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.

Feature History for Multiple Context Mode

Table 9-2 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 9-2 Feature History for Multiple Context Mode

Feature Name	Platform Releases	Feature Information
Multiple security contexts	7.0(1)	Multiple context mode was introduced. We introduced the following screens: Configuration > Context Management.
Automatic MAC address assignment	7.2(1)	Automatic assignment of MAC address to context interfaces was introduced. We modified the following screen: Configuration > Context Management > Security Contexts.
Resource management	7.2(1)	Resource management was introduced. We introduced the following screen: Configuration > Context Management > Resource Management.

Table 9-2 Feature History for Multiple Context Mode (continued)

Feature Name	Platform Releases	Feature Information
Virtual sensors for IPS	8.0(2)	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 ASA to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. We modified the following screen: Configuration > Context Management > Security Contexts.
Automatic MAC address assignment enhancements	8.0(5)/8.2(2)	The MAC address format was changed to use 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 addresses are also now persistent across 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. We modified the following screen: Configuration > Context Management > Security Contexts.
Maximum contexts increased for the ASA 5550 and 5580	8.4(1)	The maximum security contexts for the ASA 5550 was increased from 50 to 100. The maximum for the ASA 5580 was increased from 50 to 250.
Automatic MAC address assignment enabled by default	8.5(1)	Automatic MAC address assignment is now enabled by default. We modified the following screen: Configuration > Context Management > Security Contexts.

Table 9-2 Feature History for Multiple Context Mode (continued)

Feature Name	Platform Releases	Feature Information
Automatic generation of a MAC address prefix	8.6(1)	In multiple context mode, the ASA now converts the automatic MAC address generation configuration to use a default prefix. The ASA auto-generates the prefix based on the last two bytes of the interface (ASA 5500) or backplane (ASASM) MAC address. This conversion happens automatically when you reload, or if you reenable MAC address generation. The prefix method of generation provides many benefits, including a better guarantee of unique MAC addresses on a segment. If you want to change the prefix, you can reconfigure the feature with a custom prefix. The legacy method of MAC address generation is no longer available. **Note** To maintain hitless upgrade for failover pairs, the ASA does *not* convert the MAC address method in an existing configuration upon a reload if failover is enabled. However, we strongly recommend that you manually change to the prefix method of generation when using failover, especially for the ASASM. Without the prefix method, ASASMs installed in different slot numbers experience a MAC address change upon failover, and can experience traffic interruption. After upgrading, to use the prefix method of MAC address generation, reenable MAC address generation to use the default prefix. We modified the following screen: Configuration > Context Management > Security Contexts
Dynamic routing in Security Contexts	9.0(1)	EIGRP and OSPFv2 dynamic routing protocols are now supported in multiple context mode. OSPFv3, RIP, and multicast routing are not supported.
New resource type for routing table entries	9.0(1)	A new resource type, routes, was created to set the maximum number of routing table entries in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class
Site-to-Site VPN in multiple context mode	9.0(1)	Site-to-site VPN tunnels are now supported in multiple context mode.
New resource type for site-to-site VPN tunnels	9.0(1)	New resource types, vpn other and vpn burst other, were created to set the maximum number of site-to-site VPN tunnels in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class

Configuring a Cluster of ASAs

Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

This chapter includes the following sections:

- Information About ASA Clustering, page 10-1
- Licensing Requirements for ASA Clustering, page 10-21
- Prerequisites for ASA Clustering, page 10-21
- Guidelines and Limitations, page 10-23
- Default Settings, page 10-24
- Configuring ASA Clustering, page 10-24
- Managing ASA Cluster Members, page 10-50
- Monitoring the ASA Cluster, page 10-63
- Configuration Examples for ASA Clustering, page 10-66
- Feature History for ASA Clustering, page 10-71

Information About ASA Clustering

This section includes the following topics:

- How the ASA Cluster Fits into Your Network, page 10-2
- Performance Scaling Factor, page 10-2
- Cluster Members, page 10-2
- ASA Cluster Interfaces, page 10-4
- Cluster Control Link, page 10-6
- High Availability within the ASA Cluster, page 10-7
- Configuration Replication, page 10-8
- ASA Cluster Management, page 10-9
- Load Balancing Methods, page 10-10
- How the ASA Cluster Manages Connections, page 10-13

- ASA Features and Clustering, page 10-14

How the ASA Cluster Fits into Your Network

The cluster consists of up to 8 ASAs acting as a single unit. To act as a cluster, the ASAs need the following infrastructure:

- Isolated, high-speed backplane network for intra-cluster communication, known as the *cluster control link*. See the "Cluster Control Link" section on page 10-6.
- Management access to each ASA for configuration and monitoring. See the "ASA Cluster Management" section on page 10-9.

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using one of the following methods:

- Spanned EtherChannel (Recommended)—Interfaces on multiple members of the cluster are grouped into a single EtherChannel; the EtherChannel performs load balancing between units. See the "Spanned EtherChannel (Recommended)" section on page 10-10.
- Policy-Based Routing (Routed mode only)—The upstream and downstream routers perform load balancing between units using route maps and ACLs. See the "Policy-Based Routing (Routed Mode Only)" section on page 10-12.
- Equal-Cost Multi-Path Routing (Routed mode only)—The upstream and downstream routers perform load balancing between units using equal cost static or dynamic routes. See the "Equal-Cost Multi-Path Routing (Routed Mode Only)" section on page 10-13.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect a performance of approximately:

- 70% of the combined throughput
- 60% of maximum connections
- 50% of connections per second

For example, for throughput, the ASA 5585-X with SSP-40 can handle approximately 10 Gbps of real world firewall traffic when running alone. For a cluster of 8 units, the maximum combined throughput will be approximately 70% of 80 Gbps (8 units x 10 Gbps): 56 Gbps.

Cluster Members

- Bootstrap Configuration, page 10-3
- Master and Slave Unit Roles, page 10-3
- Master Unit Election, page 10-3
- New Connection Ownership, page 10-3

Bootstrap Configuration

On each device, you configure a minimal bootstrap configuration including the cluster name, cluster control link interface, and other cluster settings. The first unit on which you enable clustering typically becomes the *master* unit. When you enable clustering on subsequent units, they join the cluster as *slaves*.

Master and Slave Unit Roles

One member of the cluster is the master unit. The master unit is determined by the priority setting in the bootstrap configuration; the priority is set between 1 and 100, where 1 is the highest priority. All other members are slave units. Typically, when you first create a cluster, the first unit you add becomes the master unit simply because it is the only unit in the cluster so far.

You must perform all configuration (aside from the bootstrap configuration) on the master unit only; the configuration is then replicated to the slave units. In the case of physical assets, such as interfaces, the configuration of the master unit is mirrored on all slave units. For example, if you configure GigabitEthernet 0/1 as the inside interface and GigabitEthernet 0/0 as the outside interface, then these interfaces are also used on the slave units as inside and outside interfaces.

Some features do not scale in a cluster, and the master unit handles all traffic for those features. See the "Centralized Features" section on page 10-15.

Master Unit Election

Members of the cluster communicate over the cluster control link to elect a master unit, as follows:

1. When you enable clustering for a unit (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other units with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a unit does not receive a response from another unit with a higher priority, then it becomes master.

   Note If multiple units tie for the highest priority, the cluster unit name, and then the serial number is used to determine the master.

4. If a unit later joins the cluster with a higher priority, it does not automatically become the master unit; the existing master unit always remains as the master unless it stops responding, at which point a new master unit is elected.

Note You can manually force a unit to become the master. For centralized features, if you force a master unit change, then all connections are dropped, and you have to re-establish the connections on the new master unit. See the "Centralized Features" section on page 10-15 for a list of centralized features.

New Connection Ownership

When a new connection is directed to a member of the cluster, that unit owns both directions of the connection. If any connection packets arrive at a different unit, they are forwarded to the owner unit over the cluster control link. For best performance, proper external load balancing is required for both

directions of a flow to arrive at the same unit, and for flows to be distributed evenly between units. If a reverse flow arrives at a different unit, it is redirected back to the original unit. For more information, see the "Load Balancing Methods" section on page 10-10.

ASA Cluster Interfaces

You can configure data For example, if you have the ASA 5585-X with SSP-60interfaces as either spanned EtherChannels or as individual interfaces. All data interfaces in the cluster must be one type only.

- Interface Types, page 10-4
- Interface Type Mode, page 10-5

Interface Types

- Spanned EtherChannel

  You can group one or more interfaces per unit into an EtherChannel that spans all units in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the bridge group, not to the interface. The EtherChannel inherently provides load balancing as part of basic operation. See also the "Spanned EtherChannel (Recommended)" section on page 10-10.

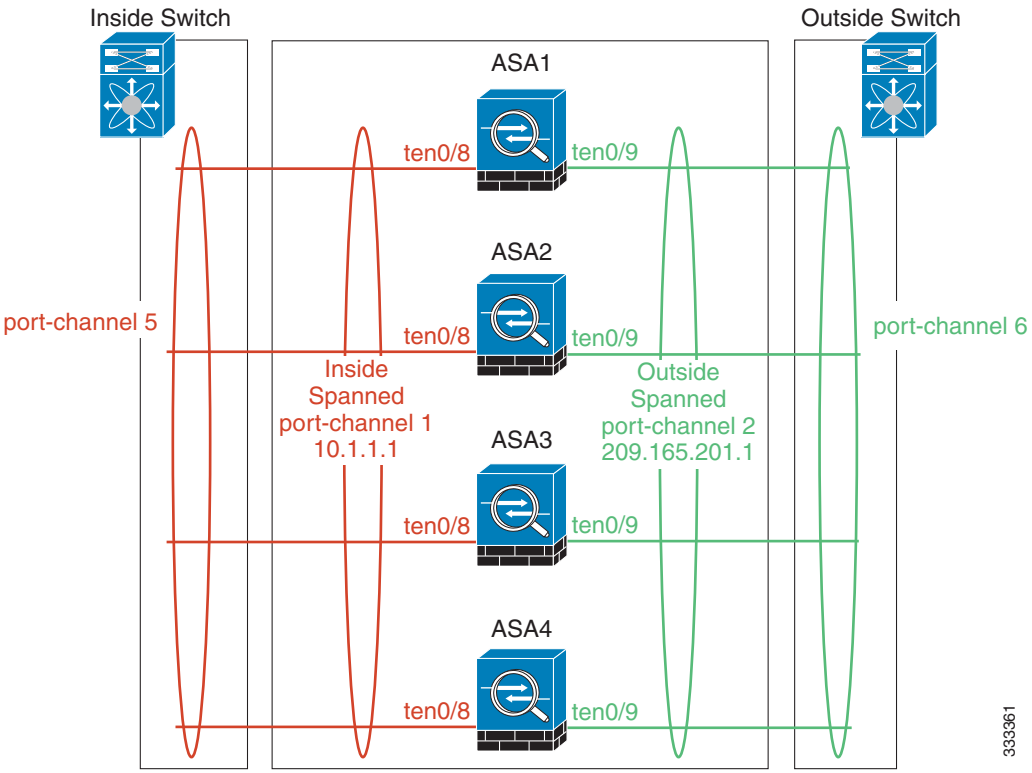

- Individual interfaces (Routed mode only)

Individual interfaces are normal routed interfaces, each with their own local IP address. Because interface configuration must be configured only on the master unit, the interface configuration lets you set a pool of IP addresses to be used for a given interface on the cluster members, including the master. The main cluster IP address is a fixed address for the cluster that always belongs to the current master unit. The main cluster IP address is a secondary IP address for the master unit; the local IP address is always the primary address for routing. The main cluster IP address provides consistent management access to an address; when a master unit changes, the main cluster IP address moves to the new master unit, so management of the cluster continues seamlessly. Load balancing, however, must be configured separately on the upstream switch in this case. For information about load balancing, see the "Load Balancing Methods" section on page 10-10.

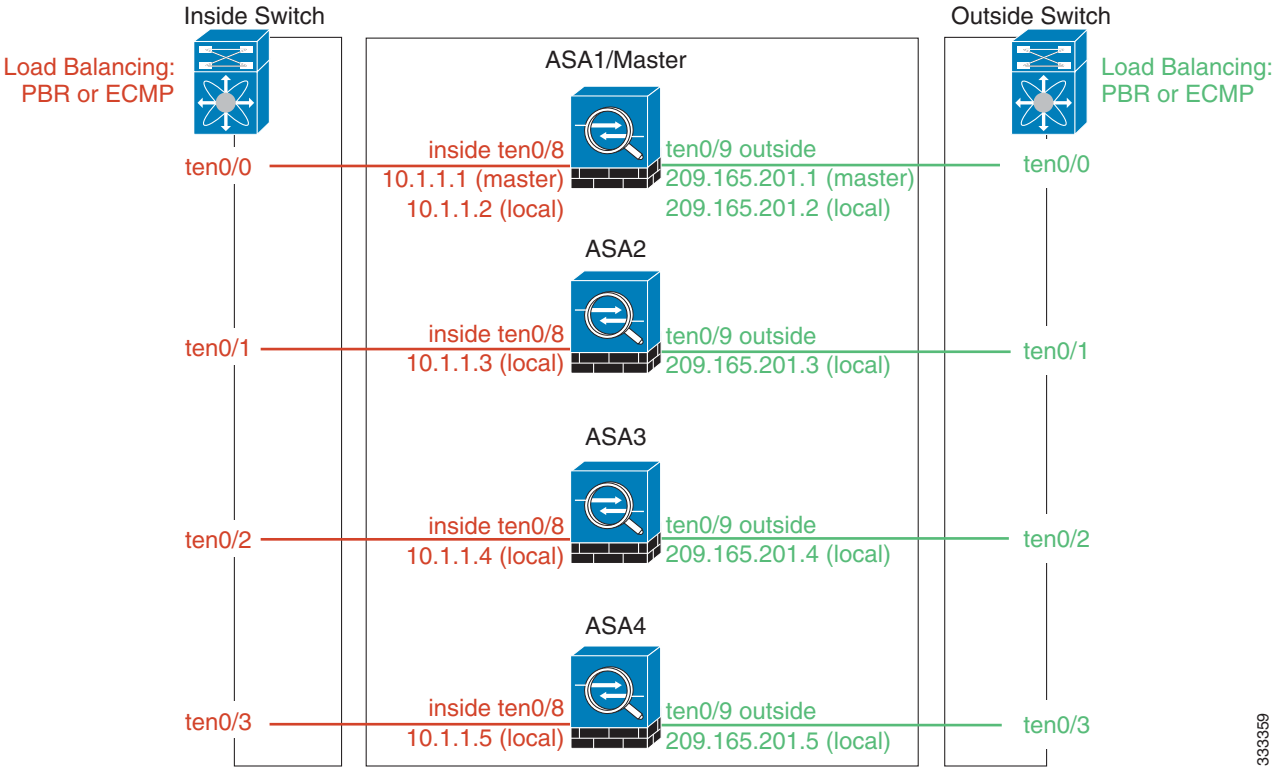

Interface Type Mode

You must choose the interface type before you configure your devices. See the following guidelines for the interface type mode:

- You can always configure the management-only interface as an individual interface (recommended), even in spanned EtherChannel mode. The management interface can be an individual interface even in transparent firewall mode.
- In spanned EtherChannel mode, if you configure the management interface as an individual interface, you cannot enable dynamic routing for the management interface. You must use a static route.
- In multiple context mode, you must choose one interface type for all contexts. For example, if you have a mix of transparent and routed mode contexts, you must use spanned EtherChannel mode for all contexts because that is the only interface type allowed for transparent mode.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link.

- Cluster Control Link Traffic Overview, page 10-6
- Cluster Control Link Network, page 10-6
- Sizing the Cluster Control Link, page 10-6
- Cluster Control Link Redundancy, page 10-7
- Cluster Control Link Failure, page 10-7

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Master election. (See the "Cluster Members" section on page 10-2.)
- Configuration replication. (See the "Configuration Replication" section on page 10-8.)
- Health monitoring. (See the "Unit Health Monitoring" section on page 10-7.)

Data traffic includes:

- State replication. (See the "Data Path Connection State Replication" section on page 10-8.)
- Connection ownership queries and data packet forwarding. (See the "Rebalancing New TCP Connections Across the Cluster" section on page 10-14.)

Cluster Control Link Network

Each cluster control link has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the ASA cluster control link interfaces.

Sizing the Cluster Control Link

You should assign an equal amount of bandwidth to the cluster control link as you assign for through traffic. We recommend that you use Ten Gigabit Ethernet interface(s) for the cluster control link. For example, if you have the ASA 5585-X with SSP-60, which can pass 14 Gbps per unit maximum in a cluster, then you should also assign interfaces to the cluster control link that can pass approximately 14 Gbps. In this case, you could use 2 Ten Gigabit Ethernet interfaces in an EtherChannel for the cluster control link, and use the rest of the interfaces as desired for data links.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. For example, state updates could consume up to 10% of the through traffic amount if through traffic consists exclusively of short-lived TCP connections. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features; for example, NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units; AAA for network access is a centralized features, so all traffic is forwarded to the master unit. When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth. A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Cluster Control Link Redundancy

If you have enough Ten Gigabit Ethernet interfaces, we recommend using an EtherChannel for the cluster control link, so you can pass traffic on multiple links in the EtherChannel while still achieving redundancy.

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS) or Virtual Port Channel (vPC) environment. All links in the EtherChannel are active. When the switch is part of a VSS or vPC, then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS or vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a spanned EtherChannel.

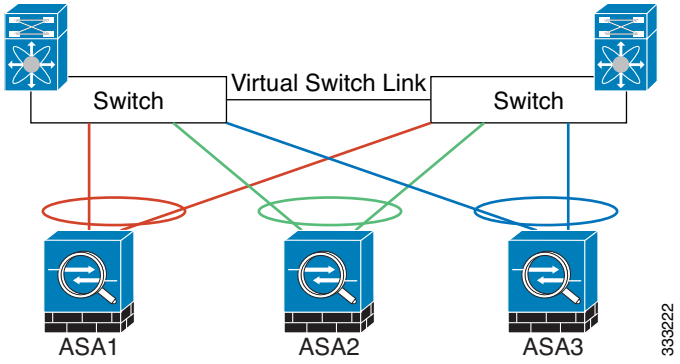

Cluster Control Link Failure

If the cluster control link line protocol goes down for a unit, then clustering is disabled; data interfaces are shut down until you manually re-enable clustering on each affected unit (after you fix the cluster control link.)

High Availability within the ASA Cluster

- Unit Health Monitoring, page 10-7
- Interface monitoring, page 10-7
- Unit Health Monitoring, page 10-7
- Data Path Connection State Replication, page 10-8

Unit Health Monitoring

The master unit monitors every slave unit by sending keepalive messages over the cluster control link periodically (the period is configurable). Each slave unit monitors the master unit using the same mechanism.

Interface monitoring

Each unit monitors the link status of all hardware interfaces in use, and reports status changes to the master unit.

- Spanned EtherChannel—Uses cluster Link Aggregation Control Protocol (cLACP). Each unit monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel. The status is reported to the master unit.
- Individual interfaces (Routed mode only)—Each unit self-monitors its interfaces and reports interface status to the master unit.

Unit or Interface Failure

When health monitoring is enabled, a unit is removed from the cluster if it fails or if its interfaces fail. If an interface fails on a particular unit, but the same interface is active on other units, then the unit is removed from the cluster.

When a unit in the cluster fails, the connections hosted by that unit are seamlessly transferred to other units; state information for traffic flows is shared over the control cluster link.

If the master unit fails, then another member of the cluster with the highest priority (lowest number) becomes the master.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure.

If the owner becomes unavailable, the first unit to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

Some traffic requires state information above the TCP or UDP layer. See Table 10-1 for clustering support or lack of support for this traffic.

Table 10-1 ASA Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	Transparent mode only.
MAC address table	Yes	Transparent mode only.
User Identity	Yes	Includes AAA rules (uauth) and identify firewall.
IPv6 Neighbor database	Yes	
Dynamic routing	Yes	
Multi-site licensing	No	
SNMP Engine ID	No	
VPN	No	VPN sessions will be disconnected if the master unit fails.

Configuration Replication

All units in the cluster share a single configuration. Except for the initial bootstrap configuration, you can only make configuration changes on the master unit, and changes are automatically replicated to all other units in the cluster.

ASA Cluster Management

- Management Network, page 10-9
- Management Interface, page 10-9
- Master Unit Management Vs. Slave Unit Management, page 10-10
- RSA Key Replication, page 10-10
- ASDM Connection Certificate IP Address Mismatch, page 10-10

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

For the management interface, we recommend using one of the dedicated management interfaces. You can configure the management interfaces as individual interfaces (for both routed and transparent modes) or as a spanned EtherChannel interface.

We recommend using individual interfaces for management, even if you use spanned EtherChannels for your data interfaces. Individual interfaces let you connect directly to each unit if necessary, while a spanned EtherChannel interface only allows remote connection to the current master unit.

Note If you use spanned EtherChannel interface mode, and configure the management interface as an individual interface, you cannot enable dynamic routing for the management interface. You must use a static route.

For an individual interface, the main cluster IP address is a fixed address for the cluster that always belongs to the current master unit. For each interface, you also configure a range of addresses so that each unit, including the current master, can use a local address from the range. The main cluster IP address provides consistent management access to an address; when a master unit changes, the main cluster IP address moves to the new master unit, so management of the cluster continues seamlessly. The local IP address is used for routing, and is also useful for troubleshooting.

For example, you can manage the cluster by connecting to the main cluster IP address, which is always attached to the current master unit. To manage an individual member, you can connect to the local IP address.

For outbound management traffic, such as TFTP or syslog, each unit uses the local IP address to connect to the server, including the master unit.

For a spanned EtherChannel interface, you can only configure one IP address, and that IP address is always attached to the master unit. You cannot connect directly to a slave unit using the EtherChannel interface; we recommend configuring the management interface as an individual interface so you can connect to each unit. You can use a device-local EtherChannel for management.

Master Unit Management Vs. Slave Unit Management

Aside from the bootstrap configuration, all management and monitoring can take place on the master unit. From the master unit, you can check runtime statistics, resource usage, or other monitoring information of all units. You can also issue a command to all units in the cluster, and replicate the console messages from slave units to the master unit.

You can monitor slave units directly if desired. Although also available from the master unit, you can perform file management on slave units (including backing up the configuration and updating images). The following functions are not available from the master unit:

- Monitoring per-unit cluster-specific statistics.
- Syslog monitoring per unit.

RSA Key Replication

When you create an RSA key on the master unit, the key is replicated to all slave units. If you have an SSH session to the main cluster IP address, you will be disconnected if the master unit fails. The new master unit uses the same key for SSH connections, so you do not need to update the cached SSH host key when you reconnect to the new master unit.

ASDM Connection Certificate IP Address Mismatch

By default, a self-signed certificate is used for the ASDM connection based on the local IP address. If you connect to the main cluster IP address using ASDM, then a warning message about a mismatched IP address appears because the certificate uses the local IP address, and not the main cluster IP address. You can ignore the message and establish the ASDM connection. However, to avoid this type of warning, you can enroll a certificate that contains the main cluster IP address and all the local IP addresses from the IP address pool. You can then use this certificate for each cluster member. For more information, see Chapter 49, “Configuring Digital Certificates.”

Load Balancing Methods

See also the “ASA Cluster Interfaces” section on page 10-4.

- Spanned EtherChannel (Recommended), page 10-10
- Policy-Based Routing (Routed Mode Only), page 10-12
- Equal-Cost Multi-Path Routing (Routed Mode Only), page 10-13

Spanned EtherChannel (Recommended)

You can group one or more interfaces per unit into an EtherChannel that spans all units in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

- Spanned EtherChannel Benefits, page 10-11
- Guidelines for Maximum Throughput, page 10-11
- Load Balancing, page 10-11
- EtherChannel Redundancy, page 10-11
- Connecting to a VSS or vPC, page 10-11

Spanned EtherChannel Benefits

The EtherChannel method of load-balancing is recommended over other methods for the following benefits:

- Faster failure discovery.
- Faster convergence time.
- Ease of configuration.

For more information about EtherChannels in general (not just for clustering), see the "EtherChannels" section on page 15-5.

Guidelines for Maximum Throughput

To achieve maximum throughput, we recommend the following:

- Use a load balancing hash algorithm that is "symmetric," meaning that packets from both directions will have the same hash, and will be sent to the same ASA in the spanned EtherChannel. We recommend using the source and destination IP address (the default) or the source and destination port as the hashing algorithm.
- Use the same type of line cards when connecting the ASAs to the switch so that hashing algorithms applied to all packets are the same.

Load Balancing

The EtherChannel link is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

The number of links in the EtherChannel affects load balancing. See the "Load Balancing" section on page 15-7 for more information.

Symmetric load balancing is not always possible. If you configure NAT, then forward and return packets will have different IP addresses and/or ports. Return traffic will be sent to a different unit based on the hash, and the cluster will have to redirect most returning traffic to the correct unit. See the "NAT" section on page 10-19 for more information.

EtherChannel Redundancy

The EtherChannel has built-in redundancy. It monitors the line protocol status of all links. If one link fails, traffic is re-balanced between remaining links. If all links in the EtherChannel fail on a particular unit, but other units are still active, then the unit is removed from the cluster.

Connecting to a VSS or vPC

You can include multiple interfaces per ASA in the spanned EtherChannel. Multiple interfaces per ASA are especially useful for connecting to both switches in a VSS or vPC. Keep in mind that an EtherChannel can have only 8 active interfaces out of 16 maximum; the remaining 8 interfaces are on standby in case of link failure. The following figure shows a 4-ASA cluster and an 8-ASA cluster, both with a total of 16 links in the EtherChannel. The active links are shown as solid lines, while the inactive links are dotted. cLACP load-balancing can automatically choose the best 8 links to be active in the EtherChannel. As shown, cLACP helps achieve load balancing at the link level.

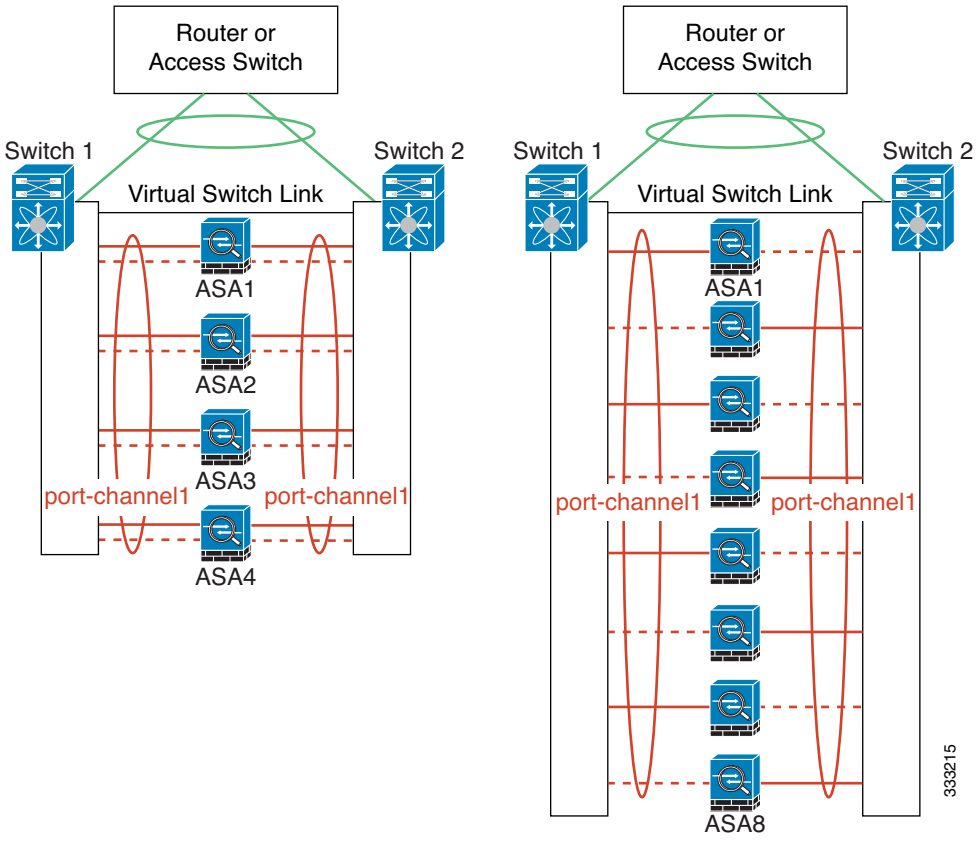

Policy-Based Routing (Routed Mode Only)

When using individual interfaces, each ASA interface maintains its own IP address and MAC address. One method of load balancing in routed mode is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure. This method might offer additional tuning options vs. spanned EtherChannel as well.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all ASAs in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same physical ASA. For example, if you have a Cisco router, redundancy can be achieved by using IOS PBR with Object Tracking. IOS Object Tracking monitors each ASA using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular ASA. See the following URLs for more details:

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml

http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtpbrtrk.html#wp1057830

Note If you use this method of load-balancing, you can use a device-local EtherChannel as an individual interface.

Equal-Cost Multi-Path Routing (Routed Mode Only)

When using individual interfaces, each ASA interface maintains its own IP address and MAC address. One method of load balancing in routed mode is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure. This method might offer additional tuning options vs. spanned EtherChannel as well.

ECMP routing can forward packets over multiple "best paths" that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then an ASA failure can cause problems; the route continues to be used, and traffic to the failed ASA will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each ASA to participate in dynamic routing.

Note If you use this method of load-balancing, you can use a device-local EtherChannel as an individual interface.

How the ASA Cluster Manages Connections

- Connection Roles, page 10-13
- Sample Data Flow, page 10-14
- Rebalancing New TCP Connections Across the Cluster, page 10-14

Connection Roles

There are 3 different ASA roles defined for each connection:

- Owner—The unit that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner.
- Director—The unit that handles owner lookup requests from forwarders and also maintains the connection state to serve as a backup if the owner fails. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and TCP ports, and sends a message to the director to register the new connection. If packets arrive at any unit other than the owner, the unit queries the director about which unit is the owner so it can forward the packets. A connection has only one director.
- Forwarder—A unit that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director (if you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required). For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Sample Data Flow

The following example shows the establishment of a new connection.

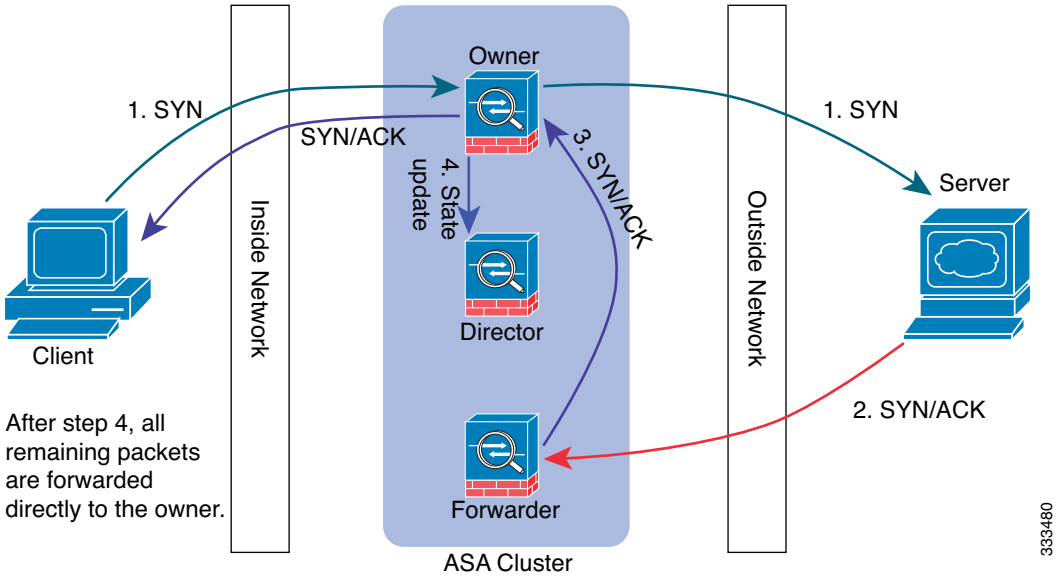

1. The SYN packet originates from the client and is delivered to an ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional units, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Rebalancing New TCP Connections Across the Cluster

If the load balancing capabilities of the upstream or downstream routers result in unbalanced flow distribution, you can configure overloaded units to redirect new TCP flows to other units. No existing flows will be moved to other units.

ASA Features and Clustering

- Unsupported Features, page 10-15

- Centralized Features, page 10-15
- Features Applied to Individual Units, page 10-16
- Dynamic Routing, page 10-17
- NAT, page 10-19
- Syslog and Netflow, page 10-20
- SNMP, page 10-20
- VPN, page 10-20
- FTP, page 10-20
- Cisco TrustSec, page 10-20

Unsupported Features

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

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

Centralized Features

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

Note Traffic for centralized features is forwarded from member units to the master unit over the cluster control link; see the "Sizing the Cluster Control Link" section on page 10-6 to ensure adequate bandwidth for the cluster control link.

If you use the rebalancing feature (see the "Rebalancing New TCP Connections Across the Cluster" section on page 10-14), traffic for centralized features may be rebalanced to non-master units before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the master unit.

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

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

Features Applied to Individual Units

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

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 8 units and with traffic evenly distributed, the conform rate actually becomes 8 times the *rate* for the cluster.
- Threat detection—Threat detection works on each unit independently; for example, the top statistics is unit-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all units, and one unit will not see all traffic.

- Resource management—Resource management in multiple context mode is enforced separately on each unit based on local usage.
- IPS module—There is no configuration sync or state sharing between IPS modules. Some IPS signatures require IPS to keep the state across multiple connections. For example, the port scanning signature is used when the IPS module detects that someone is opening many connections to one server but with different ports. In clustering, those connections will be balanced between multiple ASA devices, each of which has its own IPS module. Because these IPS modules do not share state information, the cluster may not be able to detect port scanning as a result.

Dynamic Routing

- Dynamic Routing in Spanned EtherChannel Mode, page 10-17
- Dynamic Routing in Individual Interface Mode, page 10-18

Dynamic Routing in Spanned EtherChannel Mode

In spanned EtherChannel mode, the routing process only runs on the master unit, and routes are learned through the master unit and replicated to slaves. If a routing packet arrives at a slave, it is redirected to the master unit.

Figure 10-1 Dynamic Routing in Spanned EtherChannel Mode

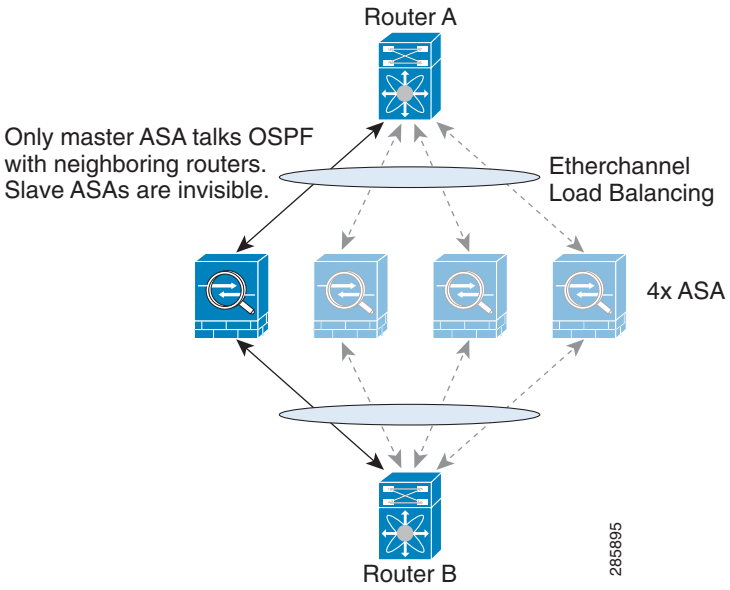

After the slave members learn the routes from the master unit, each unit makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the master unit to slave units. If there is a master unit switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster.

Dynamic Routing in Individual Interface Mode

In individual interface mode, each unit runs the routing protocol as a standalone router, and routes are learned by each unit independently.

Figure 10-2 Dynamic Routing in Individual Interface Mode

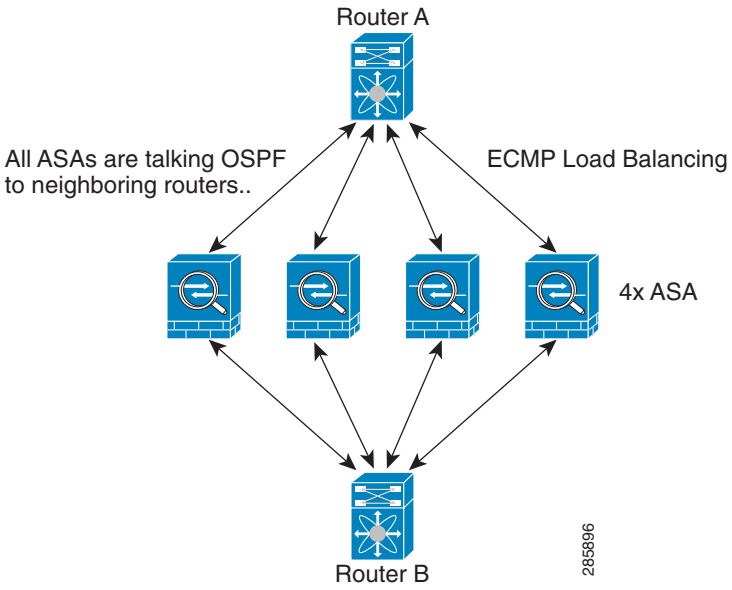

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through an ASA. ECMP is used to load balance traffic between the 4 paths. Each ASA picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each unit has a separate router ID.

Multicast Routing

- Multicast Routing in Spanned EtherChannel Mode, page 10-18
- Multicast Routing in Individual Interface Mode, page 10-18

Multicast Routing in Spanned EtherChannel Mode

In spanned EtherChannel mode, the master unit handles all multicast routing packets and data packets until fast-path forwarding is established. After the connection is established, each slave can forward multicast data packets.

Multicast Routing in Individual Interface Mode

In individual interface mode, units do not act independently with multicast. All data and routing packets are processed and forwarded by the master unit, thus avoiding packet replication.

NAT

NAT can impact the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different ASAs in the cluster because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at an ASA that is not the connection owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the main cluster IP address. This is not an issue for a spanned EtherChannel, because there is only one IP address associated with the cluster interface.
- No interface PAT on an individual interface—Interface PAT is not supported for individual interfaces.
- NAT pool address distribution—The master unit evenly pre-distributes addresses across the cluster. If a member receives a connection and they have no addresses left, the connection is dropped, even if other members still have addresses available. Make sure to include at least as many NAT addresses as there are units in the cluster to ensure that each unit receives an address. Use the **show nat pool cluster** command to see the address allocations.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- Dynamic NAT xlates managed by the master unit—The master unit maintains and replicates the xlate table to slave units. When a slave unit receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the master unit. The slave unit owns the connection.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each slave unit to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the master unit. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT. For more information about per-session PAT, see the "Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later)" section on page 39-11.

AAA for Network Access

AAA for network access consists of three components: authentication, authorization, and accounting. Authentication and accounting are implemented as centralized features on the clustering master with replication of the data structures to the cluster slaves. If a master is elected, the new master will have all the information it needs to continue uninterrupted operation of the established authenticated users and their associated authorizations. Idle and absolute timeouts for user authentications are preserved when a master unit change occurs.

Accounting is implemented as a distributed feature in a cluster. Accounting is done on a per-flow basis, so the cluster unit owning a flow will send accounting start and stop messages to the AAA server when accounting is configured for a flow.

Syslog and Netflow

- Syslog—Each unit in the cluster generates its own syslog messages. You can configure logging so that each unit uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all units in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all units look as if they come from a single unit. If you configure logging to use the local-unit name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different units. See the "Including the Device ID in Non-EMBLEM Format Syslog Messages" section on page 96-21.
- NetFlow—Each unit in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

SNMP

An SNMP agent polls each individual ASA by its local IP address. You cannot poll consolidated data for the cluster.

You should always use the local address, and not the main cluster IP address for SNMP polling. If the SNMP agent polls the main cluster IP address, if a new master is elected, the poll to the new master unit will fail.

VPN

Site-to-site VPN is a centralized feature; only the master unit supports VPN connections. Remote access VPN is not supported with clustering. VPN functionality is limited to the master unit and does not take advantage of the cluster high availability capabilities. If the master unit fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new master is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a spanned EtherChannel address, connections are automatically forwarded to the master unit. For connections to an individual interface when using PBR or ECMP, you must always connect to the main cluster IP address, not a local interface address.

VPN-related keys and certificates are replicated to all units.

FTP

If FTP data channel and control channel flows are owned by different cluster members, the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

Cisco TrustSec

Only the master unit learns security group tag (SGT) information. The master unit then populates the SGT to slaves, and slaves can make a match decision for SGT based on the security policy.

Licensing Requirements for ASA Clustering

Model	License Requirement
ASA 5580, ASA 5585-X	Cluster License. A Cluster license is required on each unit. For other feature licenses, cluster units do not require the same license on each unit. If you have feature licenses on multiple units, they combine into a single running ASA cluster license. **Note** Each unit must have the same encryption license.
All other models	No support.

Prerequisites for ASA Clustering

- Table 10-2 lists supported external hardware and software to interoperate with ASA clustering.

Table 10-2 External Hardware and Software Dependencies for ASA Clustering

Supported Hardware	Supported Software
Nexus 7000	NXOS 5.2(5)
Catalyst 6500 with Supervisor 32, 720, and 720-10GE	IOS 12.2(33)SXI7, SXI8, and SXI9

- On the switch for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the ASA to speed up the join process for new units.
- When you see slow bundling of a spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch.
- Spanned EtherChannels—For ASA *spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.

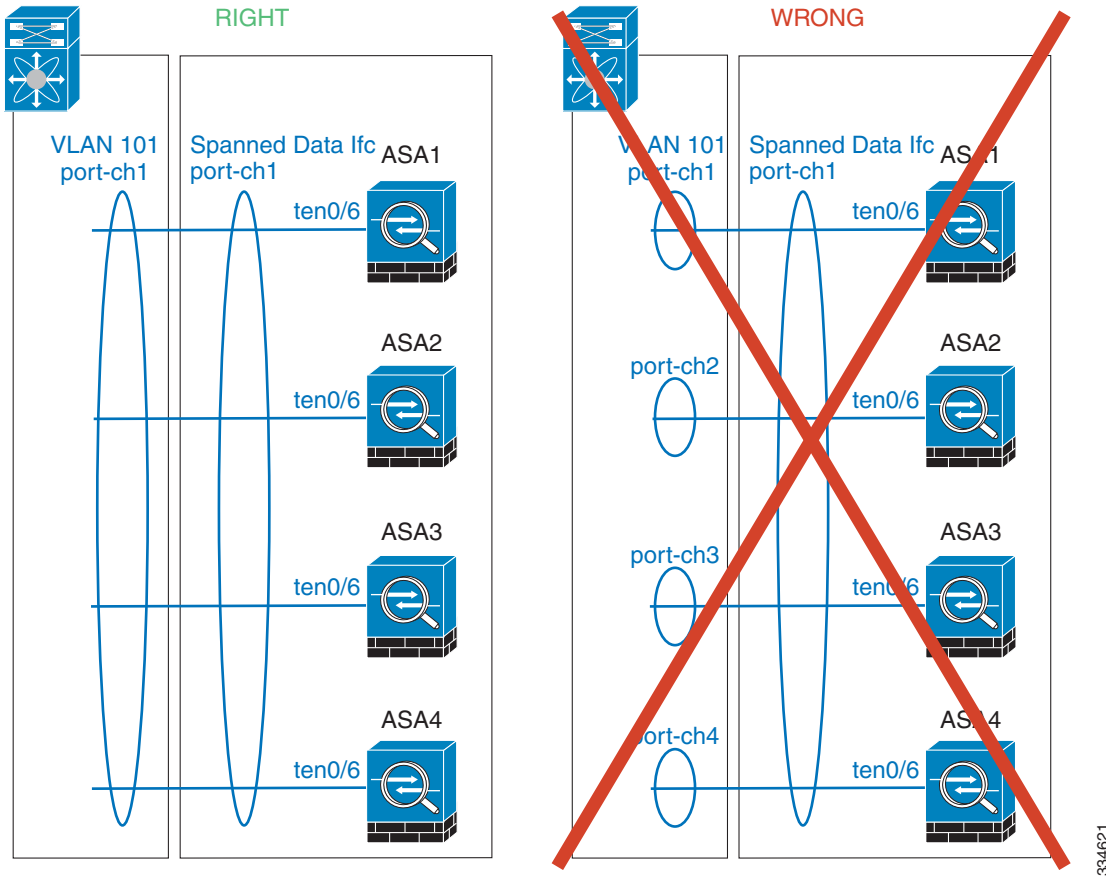

- Device-local EtherChannels—For ASA *device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple ASA EtherChannels into one EtherChannel on the switch.

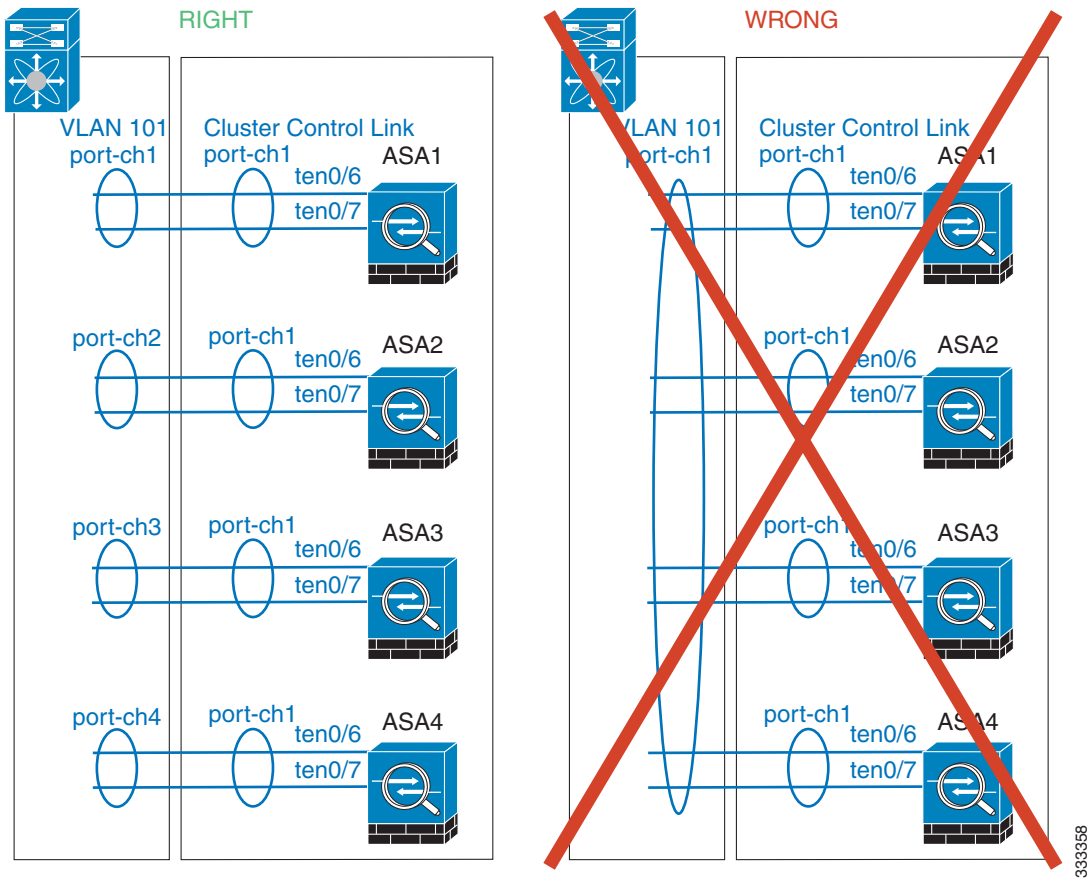

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context modes. The mode must match on each member unit.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

Failover Guidelines

Failover is not supported with clustering.

IPv6 Guidelines

Supports IPv6. However, the cluster control link is only supported using IPv4.

Model Guidelines

Supported only on the ASA 5580 and the ASA 5585-X.

For the ASA 5585-X with SSP-10 and SSP-20, which includes two Ten Gigabit Ethernet interfaces, we recommend using one interface for the cluster control link, and the other for data (you can use subinterfaces for data). Although this setup does not accommodate redundancy for the cluster control link, it does satisfy the need to size the cluster control link to match the size of the data interfaces. See the "Sizing the Cluster Control Link" section on page 10-6 for more information.

Additional Guidelines

- Each cluster can include up to 8 units.
- All units in the cluster must be the same model with the same DRAM. You do not have to have the same amount of flash memory.
- All members in a cluster must run the identical software except at the time of an image upgrade. Hitless upgrade is supported between any maintenance releases within a minor release (such as 9.0(1) to 9.0(4)), adjacent minor releases (such as 9.0 to 9.1), and last minor release of previous version to the next major release (such as 8.6 to 9.0, where 8.6 is the last version available for your model previous to 9.0).
- All members of the cluster must be in the same geographical location.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the master unit.
- When any topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.

Default Settings

- When using spanned EtherChannels, the cLACP system ID is auto-generated and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds.
- Connection rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.

Configuring ASA Clustering

- Task Flow for ASA Cluster Configuration, page 10-25
- Cabling the Cluster Units and Configuring Upstream and Downstream Equipment, page 10-25
- Configuring the Cluster Interface Mode on the Master Unit, page 10-27
- Configuring the Management Interface on Each Unit, page 10-29

- Enabling the Cluster Control Link Interface on the Master Unit, page 10-29
- Backing Up Your Configurations (Recommended), page 10-32
- Configuring Interfaces on the Master Unit, page 10-32
- Adding or Joining an ASA Cluster, page 10-44

Task Flow for ASA Cluster Configuration

To configure clustering, perform the following steps:

Step 1 Cable your equipment. Before configuring clustering, cable the cluster control link network, management network, and data networks. See the "Cabling the Cluster Units and Configuring Upstream and Downstream Equipment" section on page 10-25.

Step 2 Connect to the console port. You will need to perform configuration at the console port of each unit. See the "Accessing the Appliance Command-Line Interface" section on page 3-1. After initial configuration, you can complete your configuration using ASDM.

Step 3 Configure the interface type. You can only configure one type of interface for clustering: spanned EtherChannels or individual interfaces. See the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-27.

Step 4 Configure the management interface for each unit so you can connect to each unit using ASDM. After you configure unique IP addresses, you can use ASDM to configure the rest of your configuration. See the "Configuring the Management Interface on Each Unit" section on page 10-29.

Step 5 (Recommended) Back up each unit configuration before you enable clustering. See the "Backing Up Your Configurations (Recommended)" section on page 10-32.

Step 6 Configure interfaces for clustering on the master unit. You cannot enable clustering if the interfaces are not cluster-ready. You should complete all interface configuration before enabling clustering to avoid problems. See the "Configuring Interfaces on the Master Unit" section on page 10-32.

Step 7 Join the cluster running the High Availability and Scalability wizard. See the "Adding or Joining an ASA Cluster" section on page 10-44.

Step 8 Configure the security policy on the master unit. See the chapters in this guide to configure supported features on the master unit. The configuration is replicated to the slave units. For a list of supported and unsupported features, see the "ASA Features and Clustering" section on page 10-14.

Cabling the Cluster Units and Configuring Upstream and Downstream Equipment

Before configuring clustering, cable the cluster control link network, management network, and data networks.

Note At a minimum, an active cluster control link network is required before you configure the units to join the cluster.

You should also configure the upstream and downstream equipment. For example, if you use EtherChannels, then you should configure the upstream and downstream equipment for the EtherChannels.

Detailed Steps

Note This example uses EtherChannels for load-balancing. If you are using PBR or ECMP, your switch configuration will differ.

For example on each of 4 ASA 5585-Xs, you want to use:

- 2 Ten Gigabit Ethernet interfaces in a device-local EtherChannel for the cluster control link.
- 2 Ten Gigabit Ethernet interfaces in a spanned EtherChannel for the inside and outside network; each interface is a VLAN subinterface of the EtherChannel. Using subinterfaces lets both inside and outside interfaces take advantage of the benefits of an EtherChannel.
- 1 Management interface.

You have one switch for both the inside and outside networks.

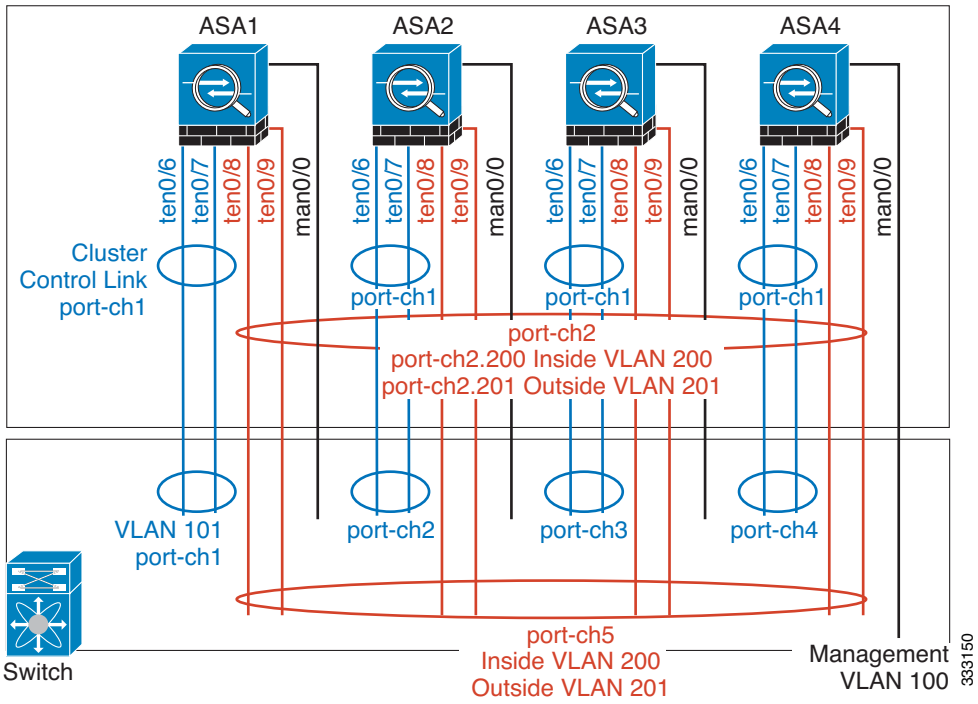

Purpose	Connect Interfaces on each of 4 ASAs	To Switch Ports
Cluster control link	TenGigabitEthernet 0/6 and TenGigabitEthernet 0/7	8 ports total For each TenGigabitEthernet 0/6 and TenGigabitEthernet 0/7 pair, configure 4 EtherChannels (1 EC for each ASA). These EtherChannels must all be on the same isolated cluster control VLAN, for example VLAN 101.
Inside and outside interfaces	TenGigabitEthernet 0/8 and TenGigabitEthernet 0/9	8 ports total Configure a single EtherChannel (across all ASAs). You can configure these VLANs and networks now or later; for example, a trunk including VLAN 200 for the inside and VLAN 201 for the outside.
Management interface	Management 0/0	4 ports total Place all interfaces on the same isolated management VLAN, for example VLAN 100.

What to Do Next

Configure the cluster interface mode from the master unit. See the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-27.

Configuring the Cluster Interface Mode on the Master Unit

You can only configure one type of interface for clustering: spanned EtherChannels or individual interfaces; you cannot mix interface types in a cluster. For exceptions for the management interface and other guidelines, see the "Interface Type Mode" section on page 10-5.

You cannot change the mode on the master unit within ASDM.

Note If you do not add slave units from the master unit, you must set the interface mode manually on all units according to this section, not just the master unit; if you add slaves from the master, ASDM sets the interface mode automatically on the slave.

Prerequisites

- Perform this procedure at the console port.
- Transparent firewall mode supports only spanned EtherChannel mode.

- For multiple context mode, configure this setting in the system execution space; you cannot configure the mode per context.

Limitations and Restrictions

When you set the cluster interface mode, your configuration cannot contain any features that are not supported with clustering. See the "Unsupported Features" section on page 10-15 for a list of unsupported features. Before you set the mode, you can check for any incompatible configuration.

Detailed Steps

	Command	Purpose
Step 1	**cluster interface-mode** {**individual** | **spanned**} **check-details** **Example:** `hostname(config)# cluster interface-mode spanned check-details`	The **check-details** command shows any incompatible configuration so you can decide whether to force the interface mode and fix your configuration later, or whether to clear the configuration when setting the interface mode; the mode is not changed with this command.
Step 2	**cluster interface-mode** {**individual** | **spanned**} [**force**] **Example:** `hostname(config)# cluster interface-mode spanned force`	Sets the interface mode for clustering. There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering. In spanned EtherChannel mode, you cannot configure a cluster IP pool for an interface. In individual interface mode, you cannot enable cluster spanning for an EtherChannel. The **force** option changes the mode without checking your configuration for incompatible settings. You need to manually fix any configuration issues after you change the mode. Because any interface configuration can only be fixed after you set the mode, we recommend using the **force** option so you can at least start from the default configuration, if present. You can re-run the **check-details** option after you set the mode for more guidance. Without the **force** option, if there is any incompatible configuration, you are prompted to clear your configuration and reload. If your configuration is compatible (rare), the mode is changed and the configuration is preserved. If you do not want to clear your configuration, you can exit the command by typing **n**. To remove the interface mode, enter the **no cluster interface-mode** command.

What to Do Next

- Configure the management interface for remote configuration. See the "Configuring the Management Interface on Each Unit" section on page 10-29.
-

Configuring the Management Interface on Each Unit

To perform configuration remotely using ASDM, configure the management interface on each unit according to this section using CLI at the console port.

Guidelines

The factory default configuration includes a pre-configured Management 0/0 interface, 192.168.1.1 (if you did not clear the configuration when you set the interface mode).

- To connect all prospective cluster members to the management network for initial configuration, you must provide a unique IP address for each unit.
- Except for the IP address used by the master unit (typically the first unit you add to the cluster), these management IP addresses are for temporary use only.
- After a slave joins the cluster, its management interface configuration is replaced by the one replicated from the master unit.

Detailed Steps

Step 1 Connect to the console port. See the "Accessing the Appliance Command-Line Interface" section on page 3-1.

Step 2 See the "Accessing ASDM Using a Non-Default Configuration (ASA 5510 and Higher)" section on page 3-9 to change the management IP address and configure ASDM access for the new address from the console port.

Step 3 Connect ASDM to the new management address. See the "Starting ASDM" section on page 3-13.

Step 4 For multiple context mode, you must also configure the cluster IP pool for the admin context management interface on the master unit only according to the "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32. You cannot later configure IP addresses within a context in the High Availability and Scalability wizard.

What to Do Next

(Optional) Configure the cluster control link. See the "Enabling the Cluster Control Link Interface on the Master Unit" section on page 10-29.

Enabling the Cluster Control Link Interface on the Master Unit

You can configure the cluster control link interface in the High Availability and Scalability wizard; if you do not use the wizard, perform the procedure in this section for the master unit; if you later add slaves from the master unit, the cluster control link interface configuration is automatically applied to the slave.

We recommend that you combine multiple cluster control link interfaces into an EtherChannel if you have enough interfaces. The EtherChannel is local to the ASA, and is not a spanned EtherChannel. We recommend that you use a Ten Gigabit Ethernet interface for the cluster control link.

Prerequisites

- Determine the size of the cluster control link by referring to the "Sizing the Cluster Control Link" section on page 10-6.
- For multiple context mode, start 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.

Restrictions

- You cannot use a VLAN subinterface as the cluster control link.
- You cannot use a Management *x*/*x* interface as the cluster control link, either alone or as an EtherChannel.
- For the ASA 5585-X with an IPS module, you cannot use the IPS module interfaces for the cluster control link.

Detailed Steps

This section includes the procedure to configure an EtherChannel. To use a single interface as the cluster control link, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23. You only need to enable the interface; do not configure a name for the interface, or any other parameters.

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 Choose **Add > EtherChannel Interface**.

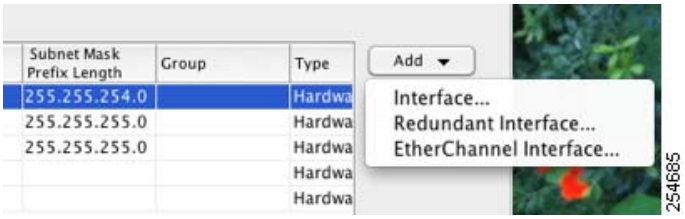

The Add EtherChannel Interface dialog box appears.

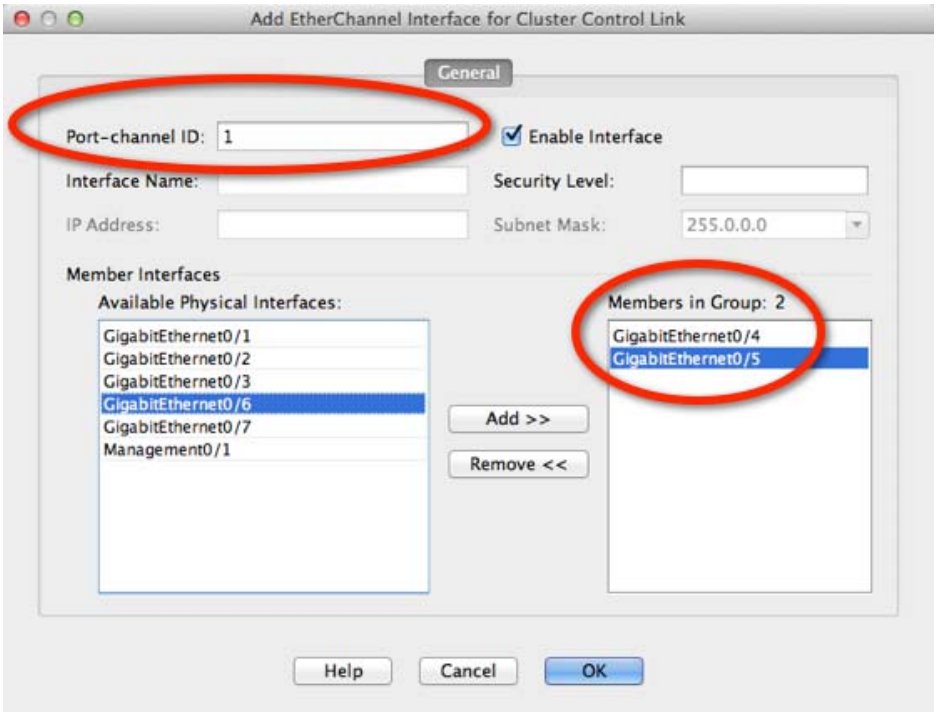

Step 3 Configure only the port channel ID and the group members. Do not configure a name for the cluster control link.

Step 4 Click **OK**.

You return to the Interfaces pane.

Step 5 Select each member interface, and click **Edit**.

The Edit Interface dialog box appears.

Step 6 Click the **Advanced** tab.

Step 7 In the EtherChannel area, from the Mode drop down list, choose **On**. We recommend using the On mode for member interfaces to reduce unnecessary traffic on the cluster control link. The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network.

Step 8 Click **OK**.

You return to the Interfaces pane.

Step 9 Click **Apply**.

What to Do Next

Back up your configuration. See the "Backing Up Your Configurations (Recommended)" section on page 10-32.

Backing Up Your Configurations (Recommended)

When you enable clustering on a slave unit, the current configuration is replaced with one synced from the master unit. If you ever want to leave the cluster entirely, it may be useful to have a backup configuration with a usable management interface configuration. See the "Leaving the Cluster" section on page 10-58 for more information.

Guidelines

Perform a backup on each unit.

Detailed Steps

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

Step 2 Back up at least the running configuration. See the "Backing Up Configurations" section on page 101-21 for a detailed procedure.

Configuring Interfaces on the Master Unit

This section describes how to configure interfaces to be compatible with clustering. You can configure data interfaces as either spanned EtherChannels or as individual interfaces. Each method uses a different load-balancing mechanism. You cannot configure both types in the same configuration. For more information, see the "ASA Cluster Interfaces" section on page 10-4.

You cannot enable clustering if the interfaces are not cluster-ready. You should complete all interface configuration before enabling clustering to avoid problems.

- Configuring Individual Interfaces (Recommended for the Management Interface), page 10-32
- Configuring Spanned EtherChannels, page 10-37

Configuring Individual Interfaces (Recommended for the Management Interface)

Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The main cluster IP address is a fixed address for the cluster that always belongs to the current master unit.

In spanned EtherChannel mode, we recommend configuring the management interface as an individual interface. Individual management interfaces let you connect directly to each unit if necessary, while a spanned EtherChannel interface only allows connection to the current master unit. See the "Management Interface" section on page 10-9 for more information.

Prerequisites

- Except for the management-only interface, you must be in individual interface mode; see the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-27.
- For multiple context mode, perform this procedure in each context. If you are not already in the context configuration mode in the Configuration > Device List pane, double-click the context name under the active device IP address.

- Individual interfaces require you to configure load balancing on neighbor devices. External load balancing is not required for the management interface. For information about load balancing, see the "Load Balancing Methods" section on page 10-10.
- (Optional) Configure the interface as a device-local EtherChannel interface, a redundant interface, and/or configure subinterfaces.
  - For an EtherChannel, see the "Configuring an EtherChannel" section on page 15-28. This EtherChannel is local to the unit, and is not a spanned EtherChannel.
  - For a redundant interface, see the "Configuring a Redundant Interface" section on page 15-25. Management interfaces cannot be redundant interfaces.
  - For subinterfaces, see the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.
- If you are connecting remotely to the management interface using ASDM, the current IP address of prospective slave units are for temporary use.
  - Each member will be assigned an IP address from the cluster IP pool defined on the master unit.
  - The cluster IP pool cannot include addresses already in use on the network, including prospective slave IP addresses.

    For example:

  a. You configure the master unit to use 10.1.1.1.
  b. Other units use 10.1.1.2, 10.1.1.3, and 10.1.1.4.
  c. When you configure the cluster IP pool on the master unit, you cannot include the .2, .3, or .4 addresses in the pool, because they are in use.
  d. Instead, you need to use other IP addresses on the network, such as .5, .6, .7, and .8.

  Note The pool needs as many addresses as there are members of the cluster, including the master unit; the original .1 address is the main cluster IP address that belongs to the current master unit.

  e. After you join the cluster, the old, temporary addresses are relinquished and can be used elsewhere.

Detailed Steps

Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

Step 2 Choose the interface row, and click **Edit**. Configure the following parameters:

a. (Required in spanned EtherChannel mode) Dedicate this interface to management only—Sets an interface to management-only mode so that it does not pass through traffic. By default, Management type interfaces are configured as management-only. In transparent mode, this command is always enabled for a Management type interface.
b. Interface Name—Enter a name up to 48 characters in length.
c. Security Level—Enter a level between 0 (lowest) and 100 (highest). See the "Security Levels" section on page 17-1 for more information.
d. Enable Interface—If the interface is not already enabled, check this check box.
e. Use Static IP—To set the IP address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported.

f. (Optional) Description—Enter a description for this interface. The description can be up to 240 characters on a single line, without carriage returns.

Note For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Step 3 To add the IPv4 cluster IP pool, and optionally a MAC address pool, click the **Advanced** tab.

g. Under ASA Cluster, create a cluster IP pool by clicking the **...** button next to the IP Address Pool field. The valid range shown is determined by the main IP address you set on the General tab.

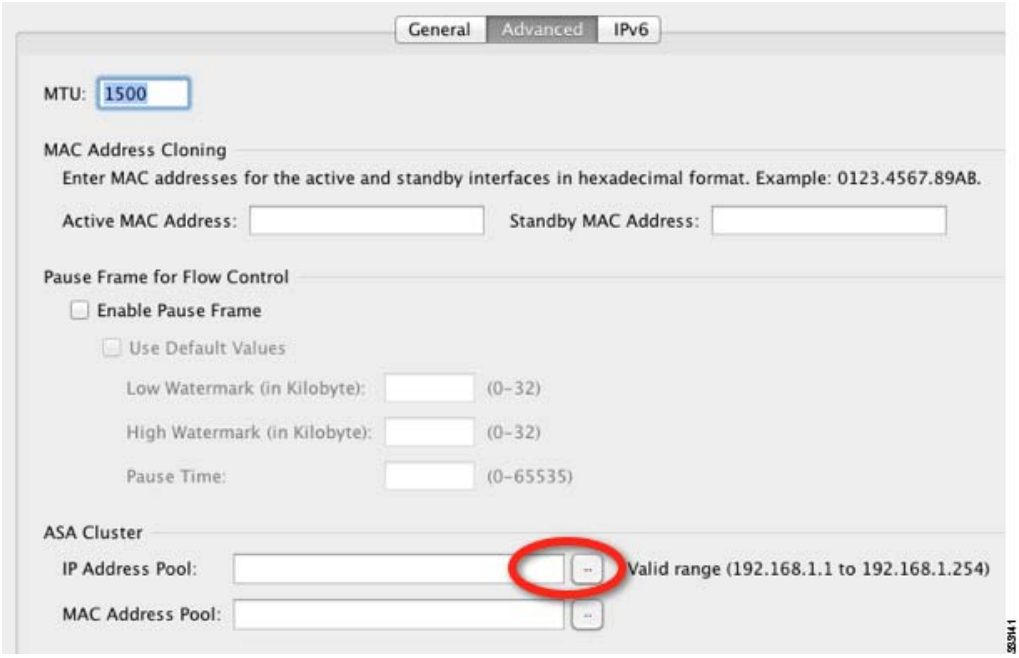

h. Click **Add**.

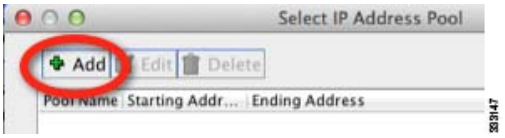

i. Configure a range of addresses that does not include the main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.

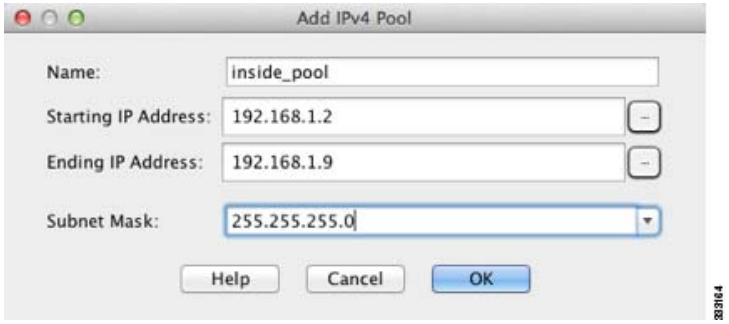

j. Click **OK** to create the new pool.

k. Select the new pool you created, and click **Assign**, and then click **OK**.

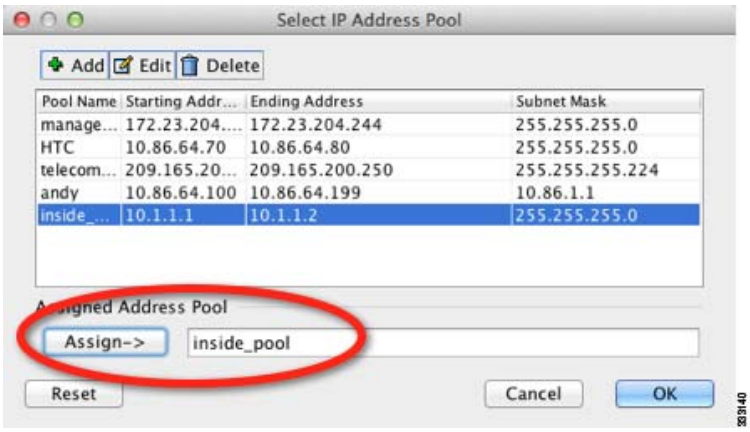

The pool name appears in the IP Address Pool field.

l. (Optional) To configure a MAC address pool, click the **...** button next to the MAC Address Pool field. Follow the screens to add a pool of MAC addresses for your interfaces. It is not common to manually configure MAC addresses for an interface, but if you have special needs to do so, then this pool is used to assign a unique MAC address to each interface.

m. For other optional parameters on the Advanced tab, see the "Configuring the MAC Address and MTU" section on page 17-13 and the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Step 4 To configure an IPv6 address, click the **IPv6** tab.

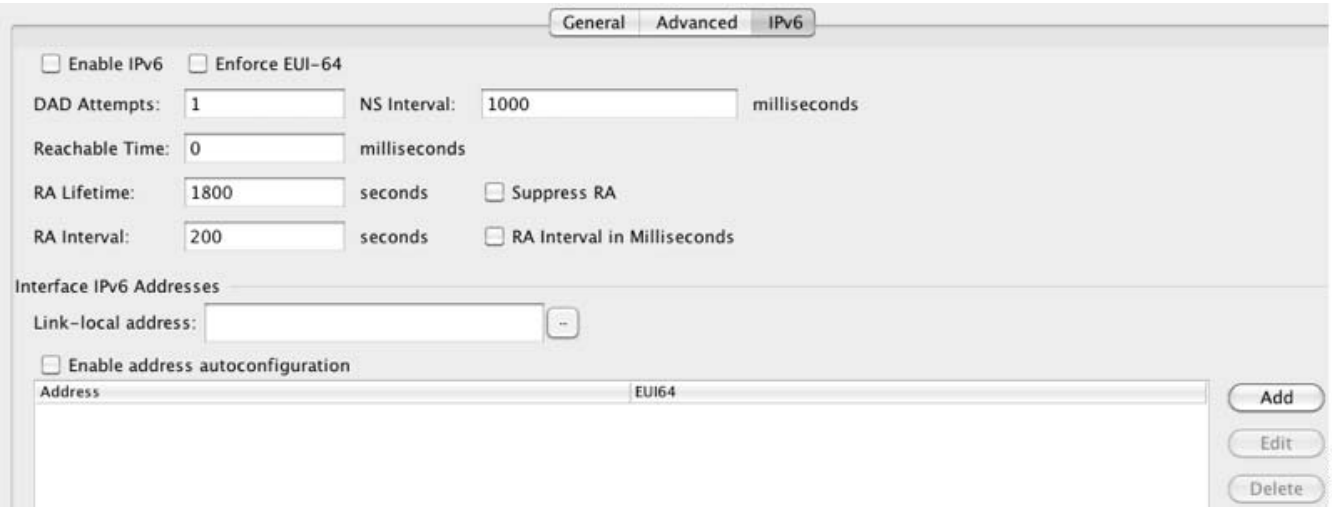

a. Check the **Enable IPv6** checkbox.

b. To manually configure a global IPv6 address, in the Interface IPv6 Addresses area, click **Add**.

   The Enable address autoconfiguration option is not supported.

   The Add IPv6 Address for Interface dialog box appears.

c. 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 A-5 for more information about IPv6 addressing.

d. Click the **...** button to configure the cluster IP pool.

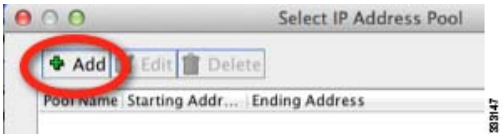

e. Click **Add**.

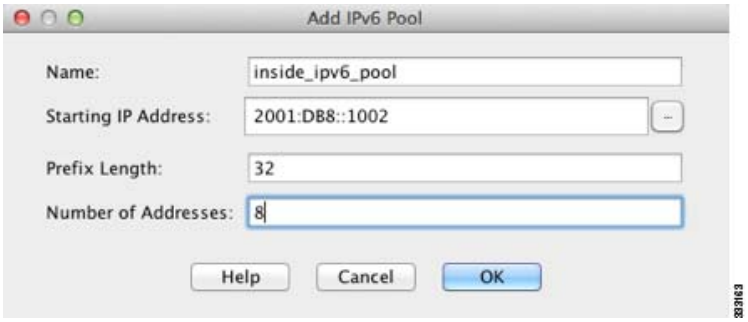

f. Configure the starting IP address (network prefix), prefix length, and number of addresses in the pool.

g. Click **OK** to create the new pool.

h. Select the new pool you created, and click **Assign**, and then click **OK**.

The pool appears in the ASA Cluster IP Pool field.

i. Click **OK**.

j. For other optional parameters on the IPv6 tab, see the "Configuring IPv6 Addressing" section on page 17-15.

Step 5 Click **OK** to return to the Interfaces pane.

Step 6 Click **Apply**.

What to Do Next

- For spanned interface mode, configure your data interfaces. See the "Configuring Spanned EtherChannels" section on page 10-37.
- For individual interface mode, join the cluster. See the "Adding or Joining an ASA Cluster" section on page 10-44.

Configuring Spanned EtherChannels

A spanned EtherChannel spans all ASAs in the cluster, and provides load balancing as part of the EtherChannel operation.

Prerequisites

- You must be in spanned EtherChannel interface mode; see the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-27.
- For multiple context mode, start 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.
- For transparent mode, configure the bridge group according to the "Configuring Bridge Groups" section on page 18-8.

Guidelines

- *Do not* specify the maximum and minimum links in the EtherChannel—We recommend that you do not specify the maximum and minimum links in the EtherChannel on either the ASA or the switch. If you need to use them, note the following:
  - The maximum links set on the ASA is the total number of active ports for the whole cluster. Be sure the maximum links value configured on the switch is not larger than the ASA value.
  - The minimum links set on the ASA is the minimum active ports to bring up a port-channel interface *per unit*. On the switch, the minimum links is the minimum links across the cluster, so this value will not match the ASA value.
- For detailed EtherChannel guidelines, limitations, and prerequisites, see the "Configuring an EtherChannel" section on page 15-28.
- See also the "EtherChannel Guidelines" section on page 15-11.

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.

Step 2 Choose **Add > EtherChannel Interface**.

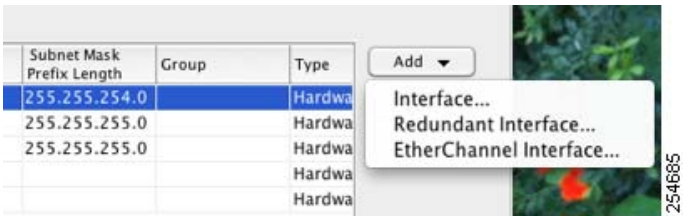

The Add EtherChannel Interface dialog box appears.

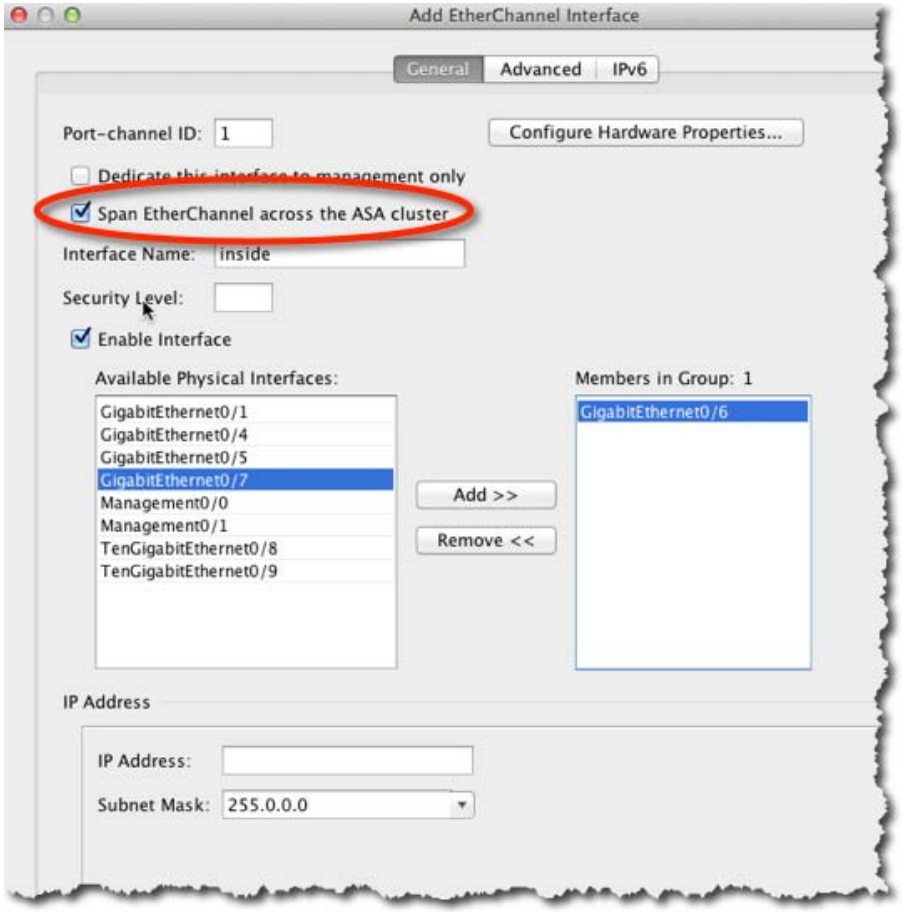

Step 3 Enable the following:

- Port Channel ID
- Span EtherChannel across the ASA cluster
- Enable Interface (checked by default)
- Members in Group—In the Members in Group box, you need to add at least one interface. Multiple interfaces in the EtherChannel per unit are useful for connecting to switches in a VSS or vPC. Keep in mind that an EtherChannel, including a spanned EtherChannel, can have only 8 active interfaces out of 16 maximum; the remaining 8 interfaces are on standby in case of link failure. For example, for a cluster of 8 ASAs, you can use a maximum of 2 interfaces on each ASA, for a total of 16 interfaces in the EtherChannel. Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

The rest of the fields on this screen are described later in this procedure.

Step 4 (Optional) To override the media type, duplex, speed, and pause frames for flow control for all member interfaces, click **Configure Hardware Properties**. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group. For more information about these settings, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Click **OK** to accept the Hardware Properties changes.

Step 5 (Optional) To configure the MAC address and optional parameters, click the **Advanced** tab.

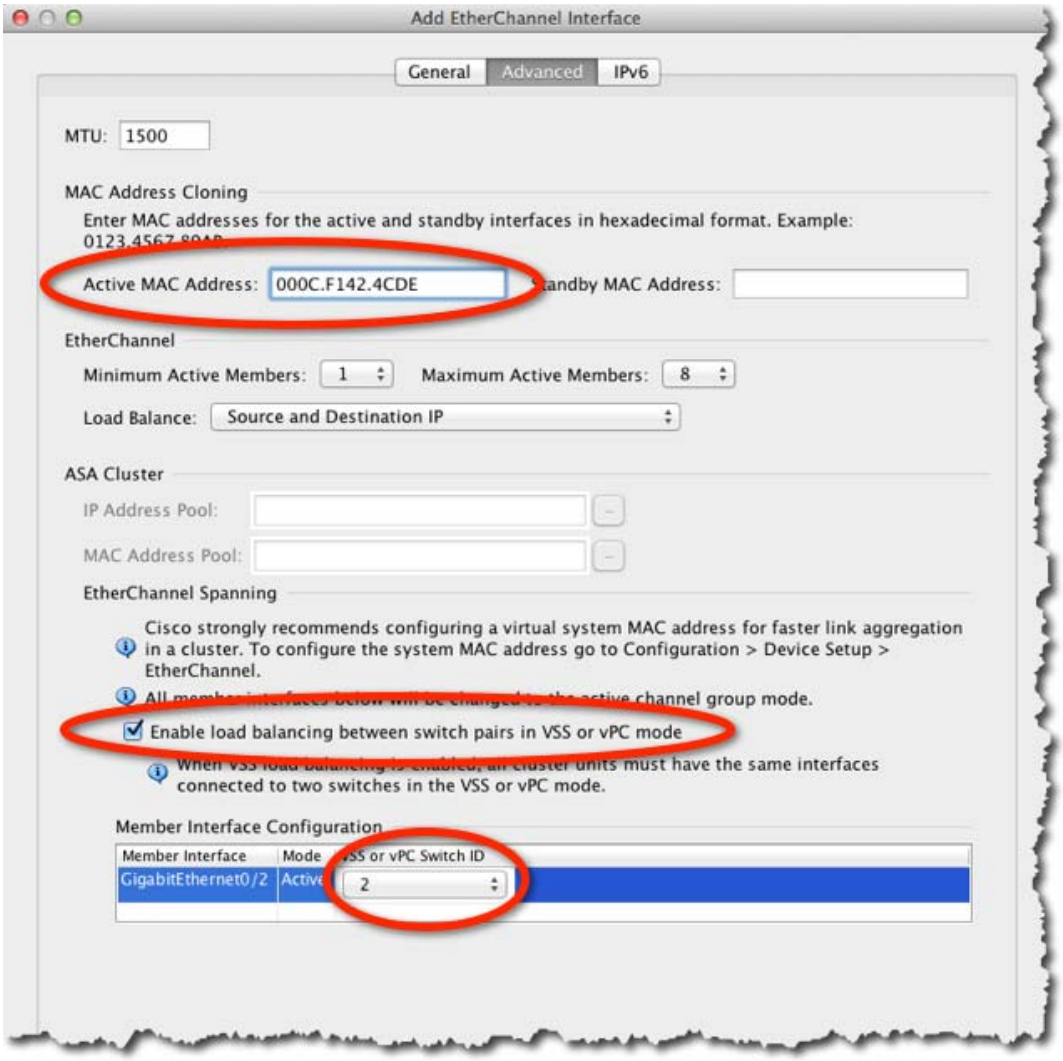

a. In single mode, under MAC Address Cloning, set a manual MAC address for the EtherChannel. Do not set the Standby MAC Address; it is ignored. You must configure a MAC address for a spanned EtherChannel so that the MAC address does not change when the current master unit leaves the cluster; with a manually-configured MAC address, the MAC address stays with the current master unit.

   In multiple context mode, if you share an interface between contexts, auto-generation of MAC addresses is enabled by default, so you only need to set the MAC address manually for a shared interface if you disable auto-generation. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

b. (Optional) If you are connecting the ASA to two switches in a VSS or vPC, then you should enable VSS load balancing by checking the **Enable load balancing between switch pairs in VSS or vPC** mode check box. This feature ensures that the physical link connections between the ASAs to the VSS (or vPC) pair are balanced.

   Under Member Interface Configuration, you must then identify to which switch a given interface is connected, 1 or 2.

c. (Optional) For information about Load Balancing, see the "Configuring an EtherChannel" section on page 15-28. For information about the MTU, see the "Configuring the MAC Address and MTU" section on page 17-13.

   We recommend that you do not set the Minimum Active Members and the Maximum Active Members. See the "Guidelines" section on page 10-38 for more information.

Step 6 (Optional) If you want to configure VLAN subinterfaces on this EtherChannel, see the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34. The rest of this procedure applies to the subinterfaces.

Step 7 (Multiple context mode) Before you complete this procedure, you need to allocate interfaces to contexts.

a. Click **OK** to accept your changes.

b. To allocate interfaces, see the "Configuring a Security Context" section on page 9-20.

c. Change to the context you want to configure: in the Device List pane, double-click the context name under the active device IP address.

d. Choose the **Configuration > Device Setup > Interfaces** pane, select the port-channel interface you want to customize, and click **Edit**.

   The Edit Interface dialog box appears.

Step 8 Click the **General** tab.

Step 9 (Transparent Mode) In the Bridge Group drop-down menu, choose the bridge group to which you want to assign this interface.

Step 10 In the Interface Name field, enter a name up to 48 characters in length.

Step 11 In the Security level field, enter a level between 0 (lowest) and 100 (highest). See the "Security Levels" section on page 17-1 for more information.

Step 12 (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

Step 13 (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.

Step 14 (Routed Mode) To configure an IPv6 address, click the **IPv6** tab.

For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

a. Check the **Enable IPv6** check box.

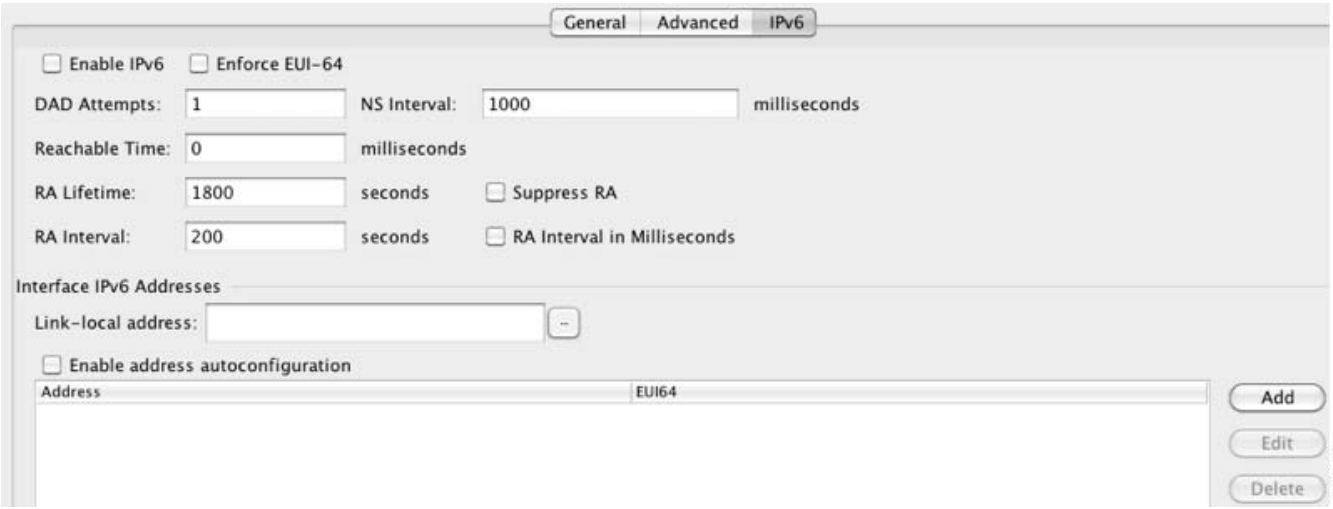

b. To manually configure a global IPv6 address, in the Interface IPv6 Addresses area, click **Add**.

The Enable address autoconfiguration option is not supported.

The Add IPv6 Address for Interface dialog box appears.

c. In the Address/Prefix Length field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:DB8::BA98:0:3210/64. See the "IPv6 Addresses" section on page A-5 for more information about IPv6 addressing.

d. (Optional) To use the Modified EUI-64 interface ID as the host address, check the **EUI-64** check box. In this case, just enter the prefix in the Address/Prefix Length field.

e. Click **OK**.

f. For other optional parameters on the IPv6 tab, see the "Configuring IPv6 Addressing" section on page 17-15.

Step 15 Click **OK** to return to the Interfaces screen.

Step 16 If a member interface is not already enabled (interfaces are disabled by default), enable it by selecting the member interface, and clicking **Edit**. Then check the **Enable Interface** check box, and click **OK**.

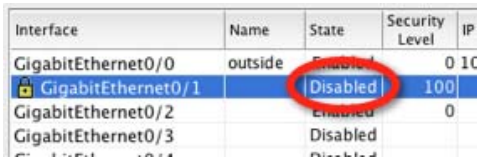

Step 17 Click **Apply**.

What to Do Next

Join the cluster. See the "Adding or Joining an ASA Cluster" section on page 10-44.

Adding or Joining an ASA Cluster

Each unit in the cluster requires a bootstrap configuration to join the cluster. Run the High Availability and Scalability wizard on one unit (that will become the master unit) to create the cluster, and then add slave units to it. If you do not want to use the wizard, or want to set advanced cluster parameters, see the "Configuring ASA Cluster Parameters" section on page 10-51.

Note For the master unit, if you want to change the default of the cLACP system ID and priority, you cannot use the wizard; you must configure the cluster manually according to the "Configuring ASA Cluster Parameters" section on page 10-51.

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.
- For multiple context mode, individual interfaces within contexts must already be cluster-ready, with a cluster IP pool assigned. You can use the wizard to configure interfaces in single mode. See the "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32.
- We recommend enabling jumbo frame reservation for use with the cluster control link. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37.
- The cluster control link interface must be in an up state.
- When you add a unit to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Detailed Steps 1—Starting the Wizard

Step 1 Choose **Wizards > High Availability and Scalability Wizard**.

Step 2 **ASA Cluster** is already selected. Click **Next**.

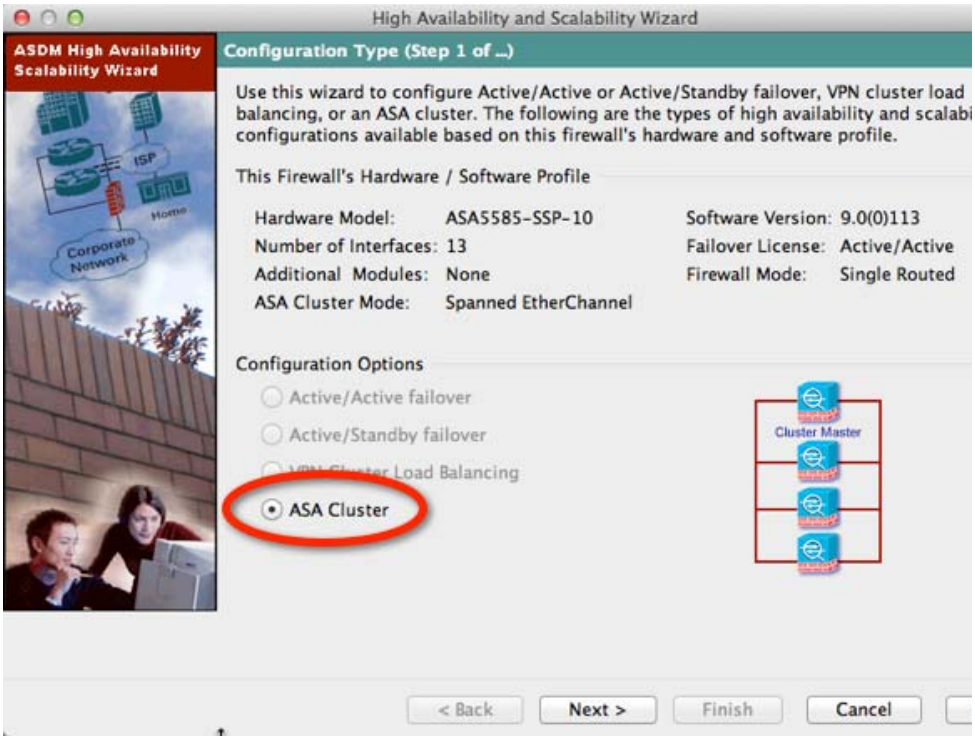

Step 3 Click **Set up a new ASA cluster**, and click **Next**.

If you click **Join an existing ASA cluster**, you add this ASA to an existing cluster.

On a master ASA in an existing cluster, the second radio button is labeled **Add another member to the cluster**.

Step 4 Click **Next**. The ASA Cluster Mode screen appears.

If you have already set the cluster interface mode (see the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-27), this screen shows the currently-configured interface mode. If you have not set the interface mode, you are prompted to exit the wizard and set the mode at the CLI before continuing.

Step 5 Click **Next**. The Interfaces screen appears.

Detailed Steps 2—Configuring Interfaces

Step 1 If you did not already configure interfaces according to the "Configuring Interfaces on the Master Unit" section on page 10-32, you can use this screen to configure some interface parameters, including changing settings for any existing interface and creating an EtherChannel for the cluster control link.

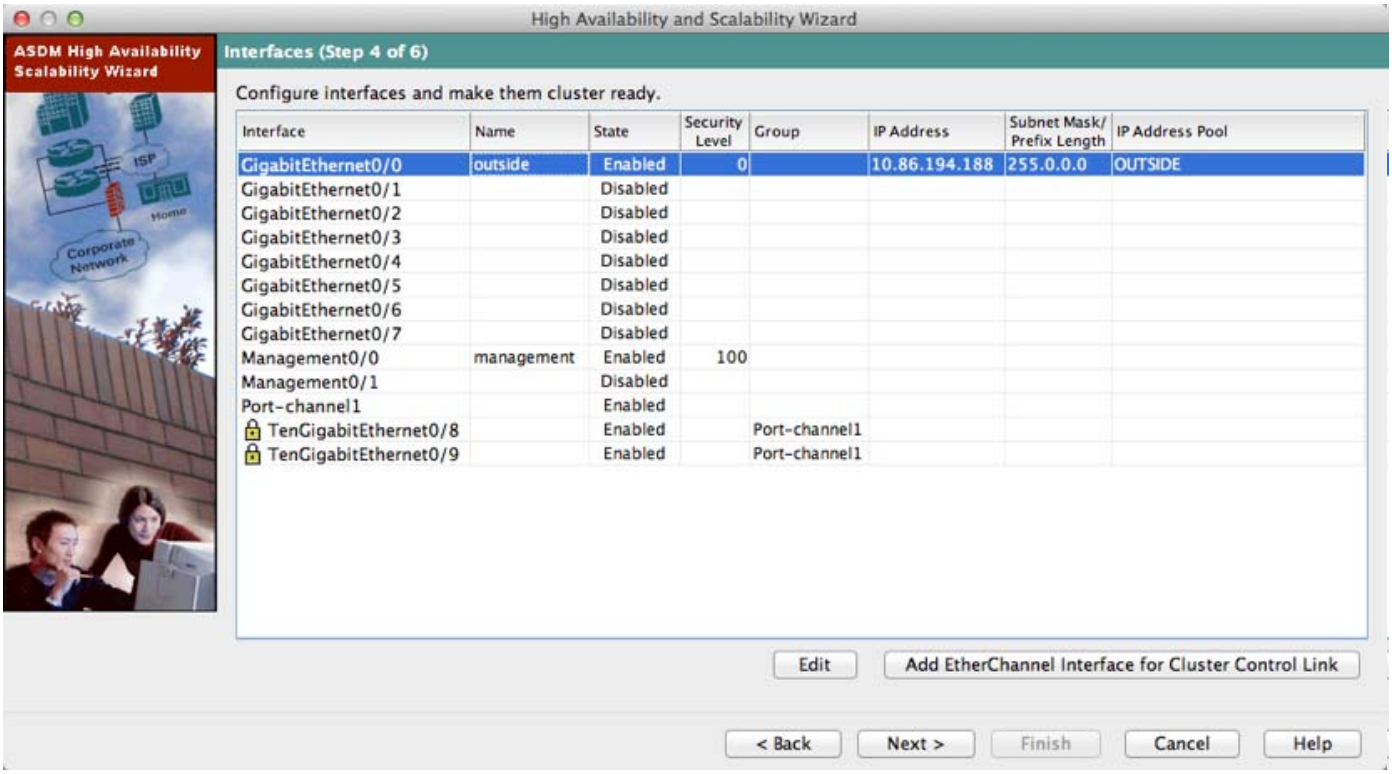

Step 2 Make any necessary interface changes; you cannot create new EtherChannels from this screen (except for the cluster control link). For more information about configuring interfaces for clustering, see the following sections:

- "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32
- "Configuring Spanned EtherChannels" section on page 10-37
- "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32

Step 3 Click **Next**.

The wizard verifies that all interfaces are cluster-ready. If you have any interfaces that are not cluster-ready, you see an error similar to the following:

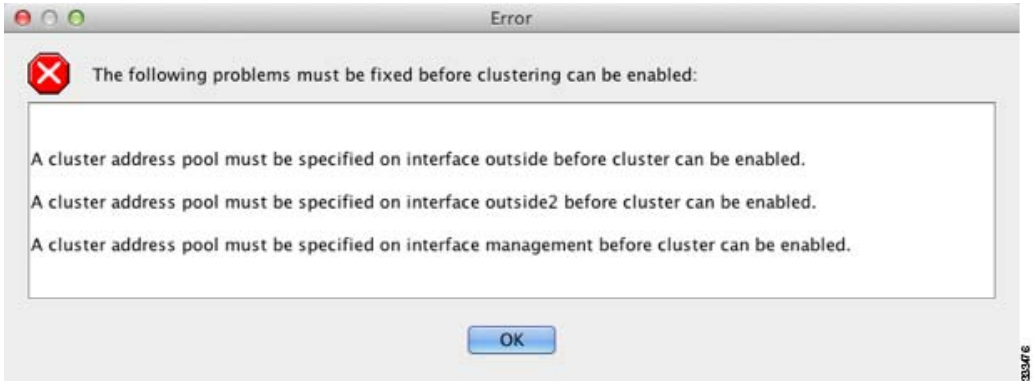

You are returned to the Interfaces screen to fix the interface configuration.

If your interface configuration passes the error checker, the ASA Cluster Configuration screen appears.

Detailed Steps 3—Configuring Bootstrap Parameters

Step 1 Configure the following bootstrap parameters:

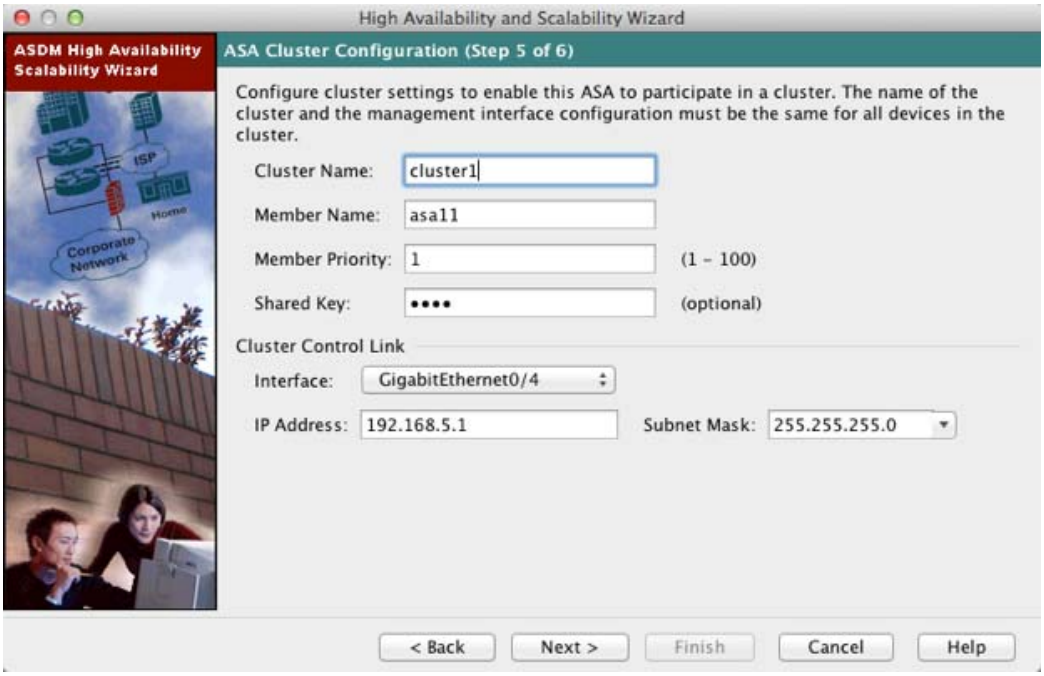

a. Cluster Name—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per unit. All members of the cluster must use the same name.

b. Member Name—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.

c. Member Priority—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.

d. (Optional) Shared Key—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.

e. Cluster Control Link—Specifies the cluster control link interface. This interface cannot have a name configured; available interfaces are shown in the drop-down list.

   – Interface—Specifies the interface ID, preferably an EtherChannel. Subinterfaces and Management type interfaces are not allowed.

   – IP Address—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.

   – Subnet Mask—Specifies the subnet mask.

Step 2 Click **Next**.

The wizard shows the cluster configuration.

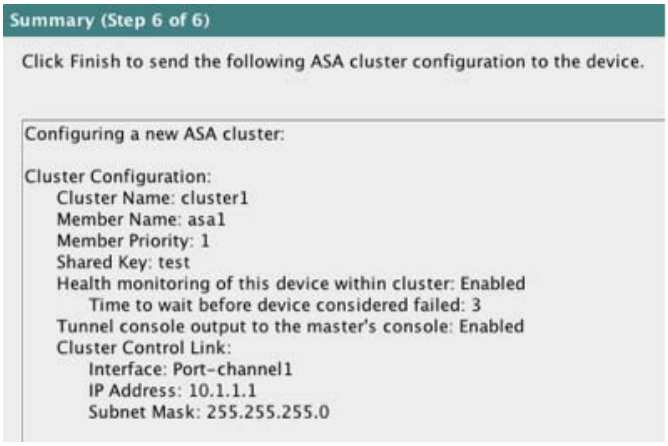

Step 3 Click **Finish**.

Step 4 The ASA scans the running configuration for incompatible commands for features that are not supported with clustering, including commands that may be present in the default configuration. Click **OK** to delete the incompatible commands. If you click **Cancel**, then clustering is not enabled.

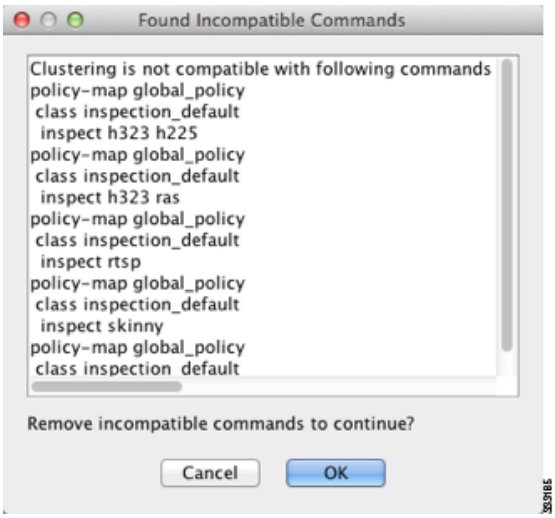

Step 5 After a period of time while ASDM enables clustering and reconnects to the ASA, the Information screen appears confirming that the ASA was added to the cluster.

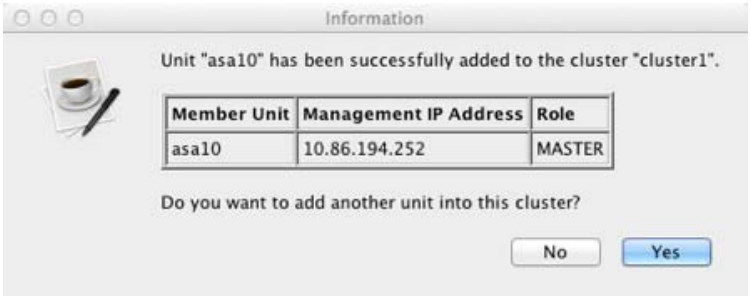

Note In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact error condition that disabled clustering; for example, the cluster control link might be down.

Detailed Steps 4—Adding Slave Units

Step 1 To add a slave unit, click **Yes**.

If you are re-running the wizard from the master, you can add slave units by choosing the **Add another member to the cluster** option when you first start the wizard.

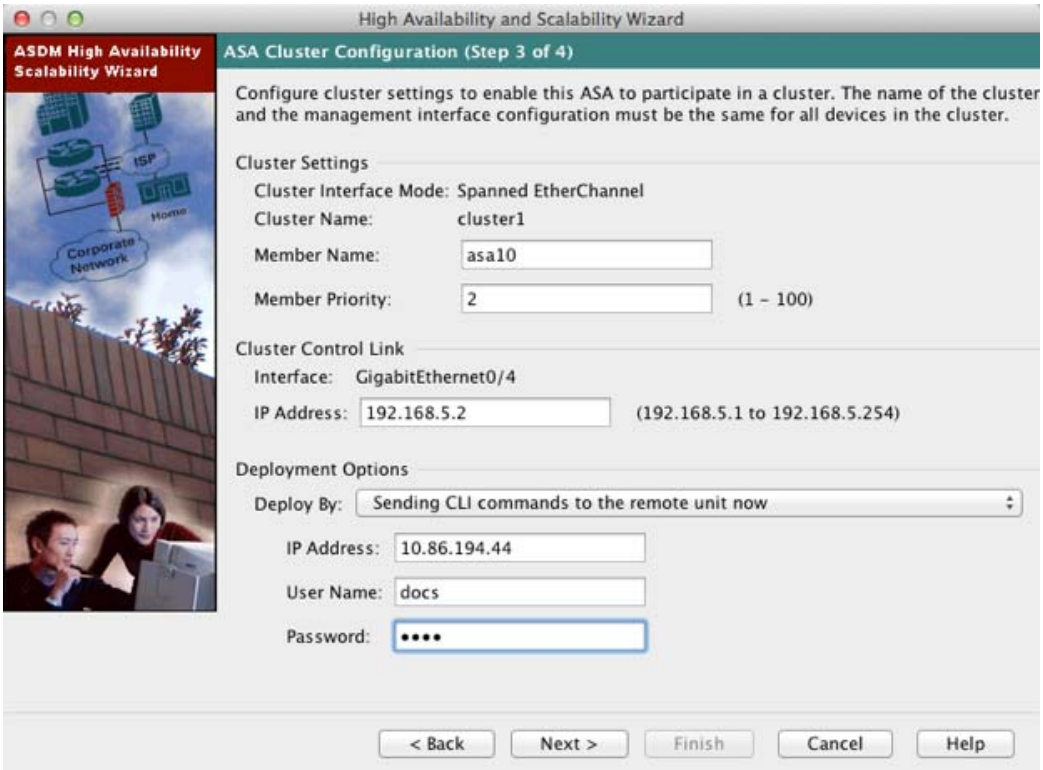

Step 2 Set the new member name, priority, and cluster control link IP address (on the same network as the other units' cluster control links).

Step 3 Under Deployment Options, choose one of the following Deploy By options:

- **Sending CLI commands to the remote unit now—**Send the bootstrap configuration to the slave (temporary) management IP address. Enter the slave management IP address, username, and password.
- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so you can cut and paste them at the slave unit CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.

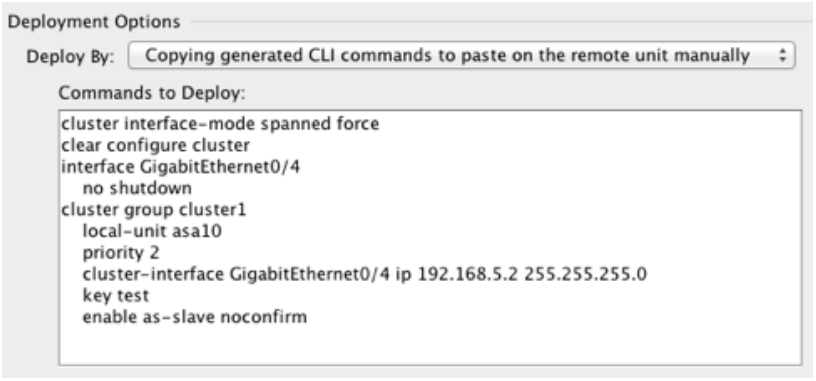

Step 4 Click **Next**. After a validation, the Summary screen appears. Click **Finish**.

The slave unit is added to the cluster.

What to Do Next

Configure the security policy on the master unit. See the chapters in this guide to configure supported features on the master unit. The configuration is replicated to the slave units. For a list of supported and unsupported features, see the "ASA Features and Clustering" section on page 10-14.

To configure advanced cluster parameters, see the "Configuring ASA Cluster Parameters" section on page 10-51.

Managing ASA Cluster Members

- Configuring ASA Cluster Parameters, page 10-51
- Adding a New Slave from the Master Unit, page 10-54
- Becoming an Inactive Member, page 10-56
- Inactivating a Slave Member from the Master Unit, page 10-57
- Leaving the Cluster, page 10-58
- Changing the Master Unit, page 10-61
- Executing a Command Cluster-Wide, page 10-62

Configuring ASA Cluster Parameters

If you do not use the wizard to add a unit to the cluster, you can configure the cluster parameters manually. If you already enabled clustering, you can edit some cluster parameters; others that cannot be edited while clustering is enabled are grayed out. This procedure also includes advanced parameters that are not included in the wizard.

Prerequisites

For multiple context mode, perform 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 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

If your device is already in the cluster, and is the master unit, then this pane is on the Cluster Configuration tab.

Step 2 Check the **Configure ASA cluster settings** check box.

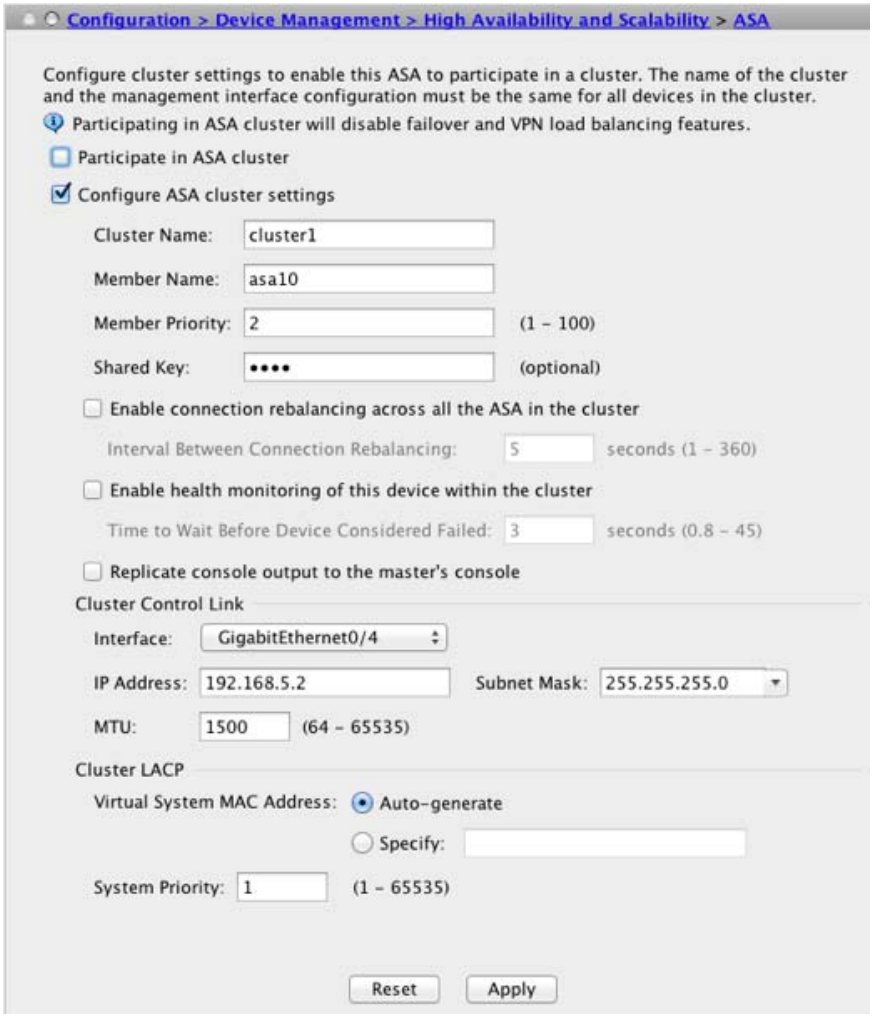

If you uncheck the check box, the settings are erased. Do not check Participate in ASA cluster until after you have set all your parameters.

Step 3 Configure the following bootstrap parameters:

a. Cluster Name—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per unit. All members of the cluster must use the same name.

b. Member Name—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.

c. Member Priority—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.

d. (Optional) Shared Key—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.

e. (Optional) Enable connection rebalancing for TCP traffic across all the ASAs in the cluster—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded

devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.

f. (Optional) Enable health monitoring of this device within the cluster—Enables the cluster health check feature, which includes unit heath monitoring and interface health monitoring. **Note**: When you are adding new units to the cluster, and making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete. You can re-enable this feature after cluster and topology changes are complete. Keepalive messages between members determine member health. If a unit does not receive any keepalive message from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead. Interface status messages detect link failure. Health check is enabled by default, and we recommend keeping it enabled. The holdime determines the amount of time between keepalive messages, between .8 and 45 seconds. By default, the holdtime is 3 seconds. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.

Note When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you must disable the health check. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check.

g. (Optional) Replicate console output to the master's console—Enables console replication from slave units to the master unit. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, slave units send the console messages to the master unit so you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.

h. Cluster Control Link—Specifies the cluster control link interface. This interface cannot have a name configured; available interfaces are shown in the drop-down list.

- Interface—Specifies the interface ID, preferably an EtherChannel. Subinterfaces and Management type interfaces are not allowed.
- IP Address—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.
- Subnet Mask—Specifies the subnet mask.
- (Optional) MTU—Specifies the maximum transmission unit for the cluster control link interface, between 64 and 65,535 bytes. Data that is larger than the MTU value is fragmented before being sent. The default MTU is 1500 bytes. We suggest stting the MTU to 1600 bytes or greater, which requires you to enable jumbo frame reservation. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37.

i. (Optional) Cluster LACP—When using spanned EtherChannels, the ASA uses cLACP to negotiate the EtherChannel with the neighbor switch. ASAs in a cluster collaborate in cLACP negotiation so that they appear as a single (virtual) device to the switch.

- Virtual System MAC Address—Sets the cLACP system ID, which is in the format of a MAC address. All ASAs use the same system ID: auto-generated by the master unit (the default) and replicated to all slaves; or manually specified in the form *H.H.H*, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units. However, you cannot change this value after you enable clustering.

– System Priority—Sets the system priority, between 1 and 65535. The priority is used to decide which unit is in charge of making a bundling decision. By default, the ASA uses priority 1, which is the highest priority. The priority needs to be higher than the priority on the switch. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units. However, you cannot change this value after you enable clustering.

Step 4 Check the **Participate in ASA cluster** check box to join the cluster.

Step 5 Click **Apply**.

Adding a New Slave from the Master Unit

You can add additional slaves to the cluster from the master unit. You can also add slaves using the High Availability and Scalability wizard. Adding a slave from the master unit has the benefit of configuring the cluster control link and setting the cluster interface mode on each slave unit you add.

You can alternatively log into the slave unit and configure clustering on the unit according to the "Configuring ASA Cluster Parameters" section on page 10-51. However, after you enable clustering, your ASDM session will be disconnected.

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.
- If you want to send the bootstrap configuration over the management network, be sure the slave unit has an accessible IP address. See the "Configuring the Management Interface on Each Unit" section on page 10-29.

Detailed Steps

Step 1 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Members**.

Step 2 Click **Add**.

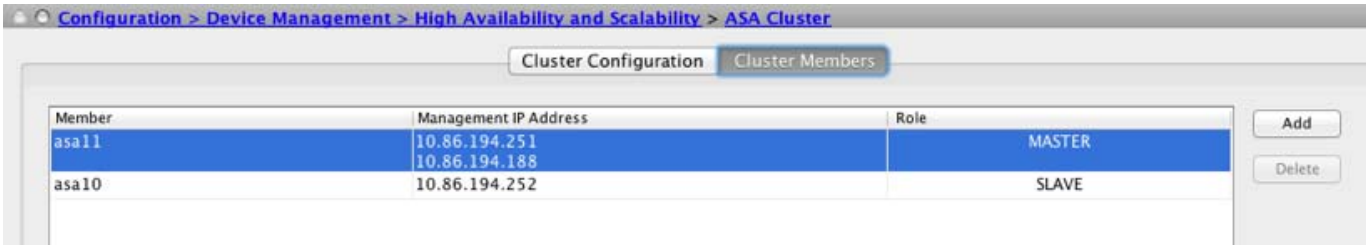

Step 3 Configure the following parameters:

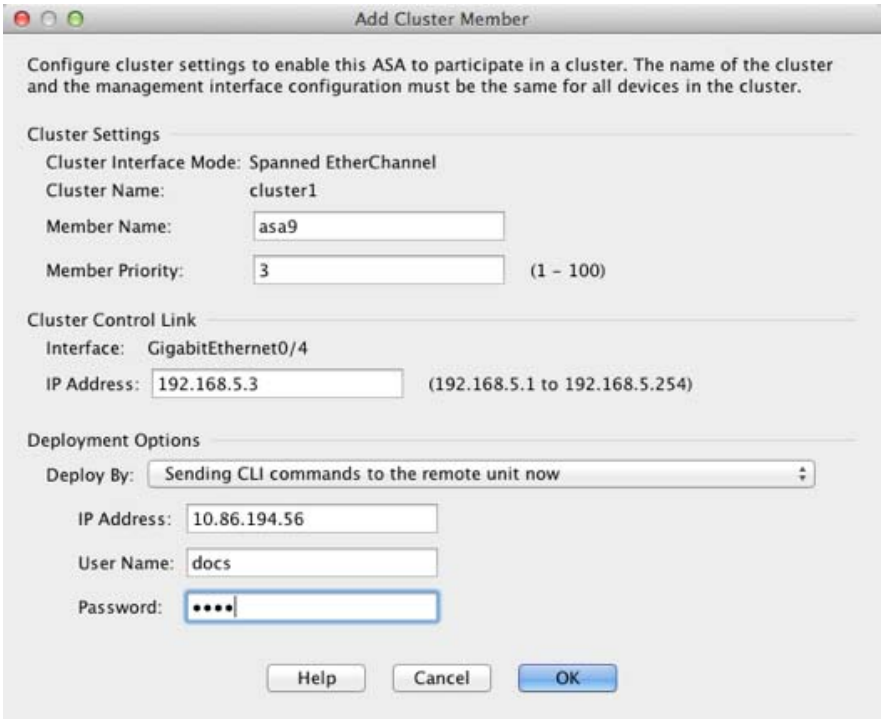

a. Member Name—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.

b. Member Priority—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.

c. Cluster Control Link > IP Address—Specifies a unique IP address for this member for the cluster control link, on the same network as the master cluster control link.

d. Under Deployment Options, choose one of the following Deploy By options:

   – **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the slave (temporary) management IP address. Enter the slave management IP address, username, and password.

   – **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so you can cut and paste them at the slave unit CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.

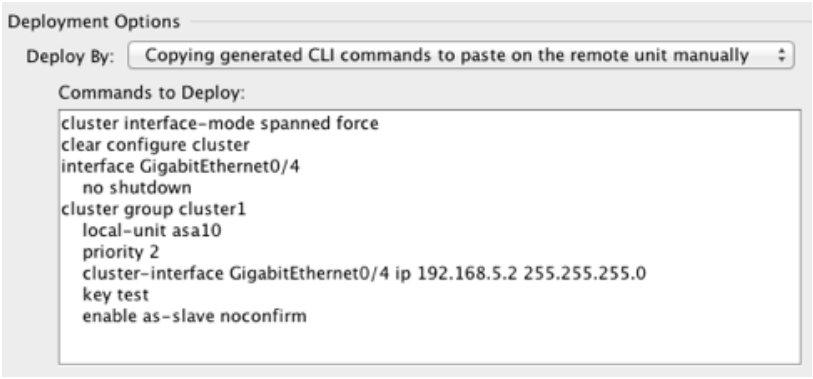

Step 4 Click **OK**, then **Apply**.

Becoming an Inactive Member

To become an inactive member of the cluster, disable clustering on the unit while leaving the clustering configuration intact.

Note When the ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the unit altogether from the cluster. See the "Leaving the Cluster" section on page 10-58. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, the management interface is not accessible (because it uses the same IP address as the master unit). You must use the console port for any further configuration.

Prerequisites

- For multiple context mode, perform 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 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

If your device is already in the cluster, and is the master unit, then this pane is on the Cluster Configuration tab.

Step 2 Uncheck the **Participate in ASA cluster** check box.

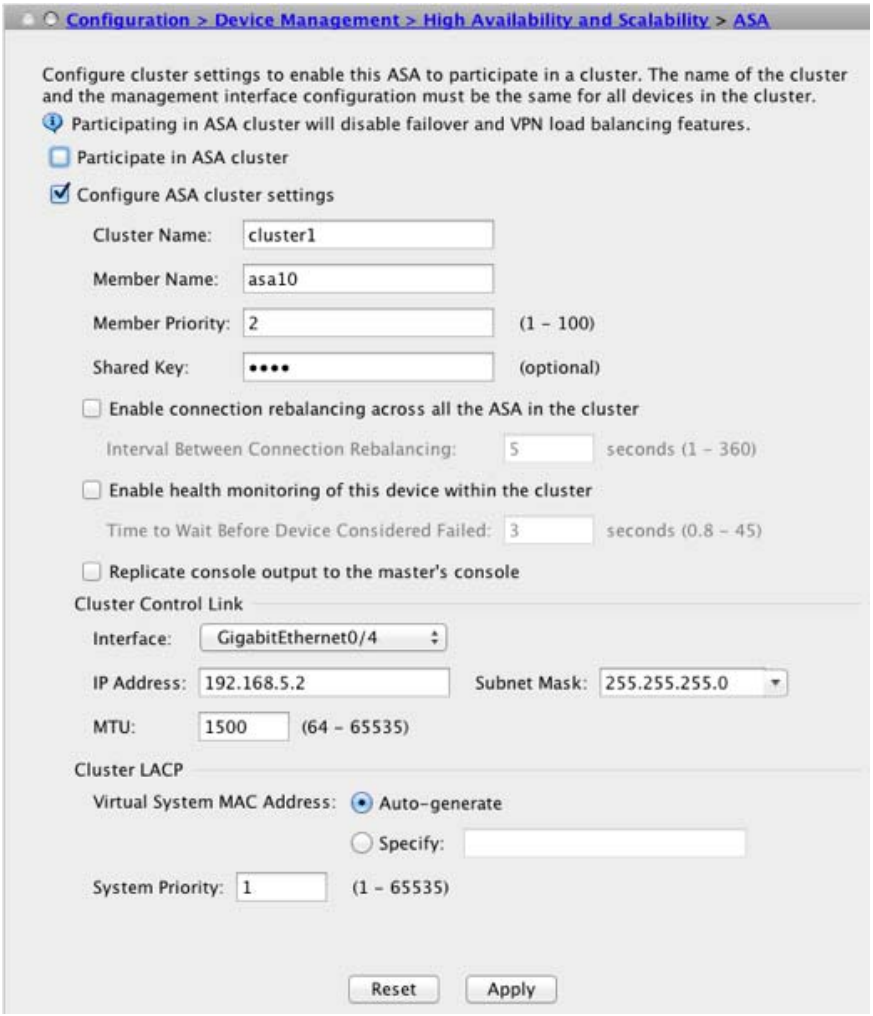

Step 3 Click **Apply**.

If this unit was the master unit, a new master election takes place, and a different member becomes the master unit.

The cluster configuration is maintained, so you can enable clustering again later.

Inactivating a Slave Member from the Master Unit

To inactivate a slave member, perform the following steps.

Note When the ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the unit altogether from the cluster. See the "Leaving the Cluster" section on page 10-58. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, the management interface is not accessible (because it uses the same IP address as the master unit). You must use the console port for any further configuration.

Prerequisites

For multiple context mode, perform 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 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

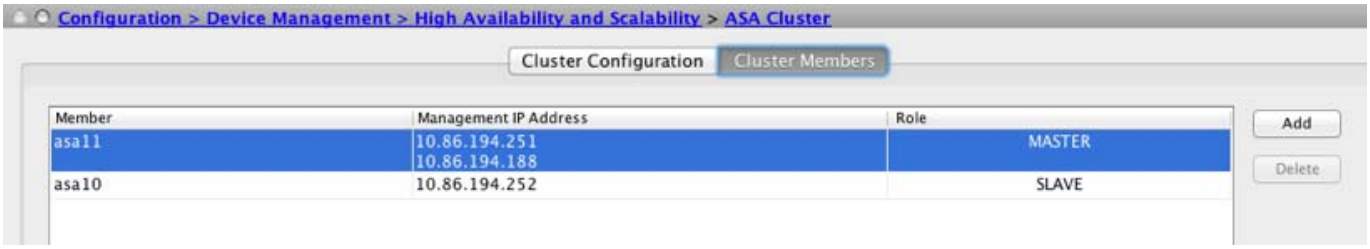

Step 2 Select the slave you want to remove, and click **Delete**.

The slave bootstrap configuration remains intact, so you can later re-add the slave without losing your configuration.

Step 3 Click **Apply**.

Leaving the Cluster

If you want to leave the cluster altogether, you need to remove the entire cluster bootstrap configuration. In practice, you need to clear or replace your entire configuration.

Prerequisites

Because the current configuration on each member is the same (synced from the master unit), leaving the cluster also means either restoring a pre-clustering configuration from backup, or clearing your configuration and starting over.

- Using ASDM, you can restore a backup running configuration. However, you also need to use the CLI at the console port to disable the cluster interface mode. See the "Backing Up Your Configurations (Recommended)" section on page 10-32 for more information about backing up the configuration before joining the cluster.
- If you do not have a backup configuration, you need to use the CLI to reconfigure the ASA for ASDM management access.
- Using CLI, you must use the console port; you cannot enable or disable clustering from a remote CLI connection.

Detailed Steps—CLI from the Console Port

	Command	Purpose

Step 1	For a slave unit:
	cluster group *cluster_name*
	no enable
	
	Example:
	`hostname(config)# cluster group cluster1`
	`hostname(cfg-cluster)# no enable`
		Disables clustering. You cannot make configuration changes while clustering is enabled on a slave unit.

Step 2	**clear configure all** **Example:** `hostname(config)# clear configure all`	Clears the running configuration.

Step 3	`no cluster interface-mode`
	
	Example:
	`hostname(config)# no cluster interface-mode`
		Disables cluster interface mode. The mode is not stored in the configuration and must be reset manually.

Step 4	If you have a backup configuration:
	copy *backup_cfg* **running-config**
	
	Example:
	`hostname(config)# copy backup_cluster.cfg running-config`
	
	`Source filename [backup_cluster.cfg]?`
	
	`Destination filename [running-config]?`
	`hostname(config)#`
		Copies the backup configuration to the running configuration.

Step 5	`write memory`
	
	Example:
	`hostname(config)# write memory`
		Saves the configuration to startup.

Step 6	If you do not have a backup configuration, reconfigure management access according to Chapter 3, "Getting Started."	

Detailed Steps—ASDM

Step 1 Connect ASDM directly to the cluster member unit using the individual management IP address, not the main cluster IP address.

To view the member IP addresses from the master unit, choose **Home > Cluster Dashboard**, and view the Management IP Address column.

Step 2 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Configuration**.

If your device is already in the cluster, and is the master unit, then this pane is on the Cluster Configuration tab.

Step 3 Uncheck the **Participate in ASA cluster** check box.

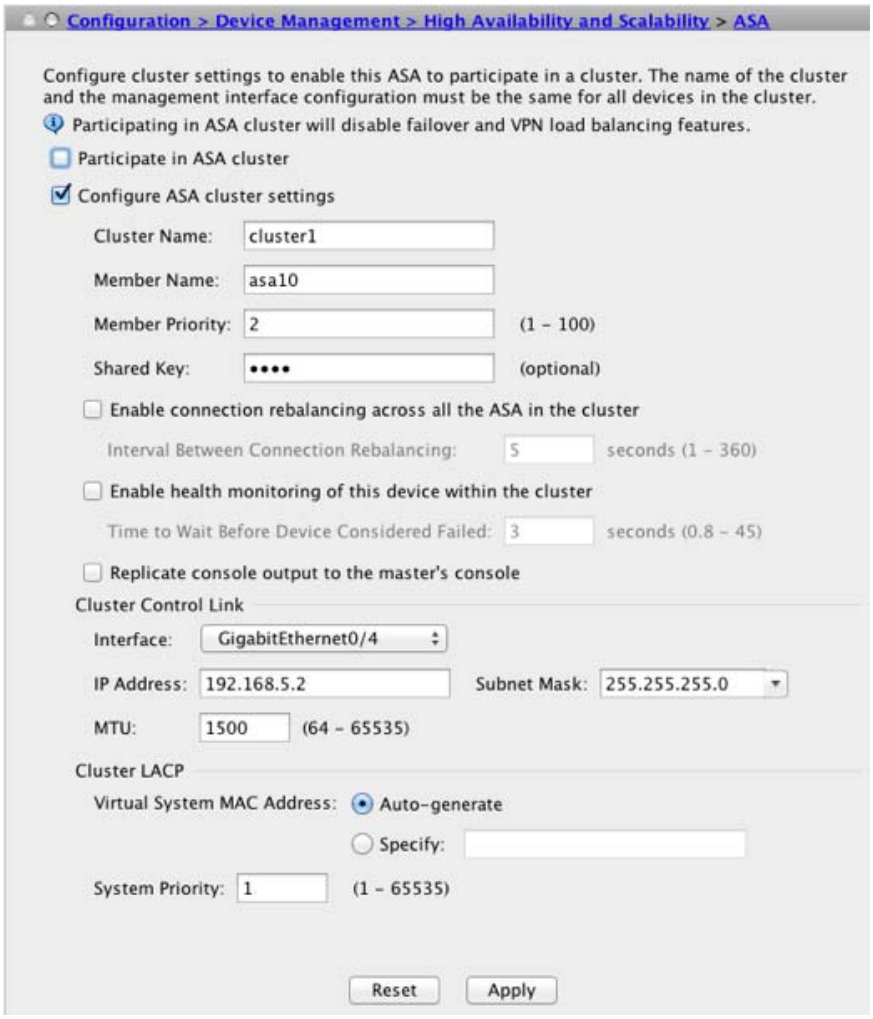

Step 4 Click **Apply**.

If this unit was the master unit, a new master election takes place, and a different member becomes the master unit.

Step 5 You are prompted to exit ASDM.

Step 6 Connect to the ASA console port, and enter the following command:

```
hostname(config)# no cluster interface-mode
```

Step 7 Reconnect ASDM to the same management IP address.

Step 8 Choose **Tools > Restore Configurations**.

The Restore Configurations dialog box appears.

Step 9 Browse to the backup zip file (that you created as part of the "Backing Up Your Configurations (Recommended)" section on page 10-32), and click **Next**.

Step 10 Check at a minimum the **Running Configuration** check box, and click **Restore**.

The Running Configuration Restore dialog box appears asking if you want to replace or merge with the current configuration.

Step 11 Click **Replace**.

Step 12 When completed, exit the restore dialog box, and exit ASDM.

Step 13 Reconnect ASDM to the *original* management IP in the backup running configuration.

Changing the Master Unit

To change the master unit, perform the following steps.

Note For centralized features, if you force a master unit change, then all connections are dropped, and you have to re-establish the connections on the new master unit. See the "Centralized Features" section on page 10-15 for a list of centralized features.

Prerequisites

For multiple context mode, perform 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 Choose **Monitoring > ASA Cluster > Cluster Summary**.

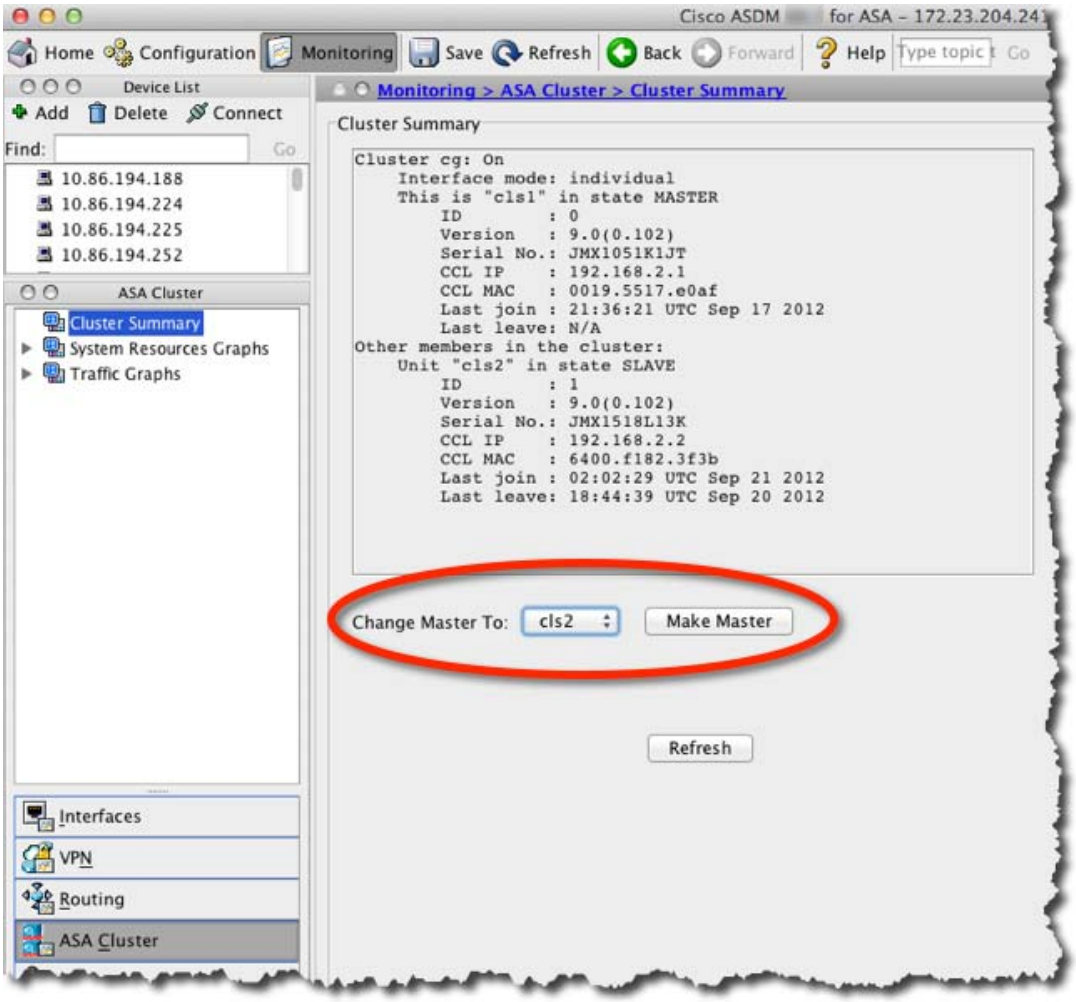

Step 2 In the Change Master To drop-down menu, choose a slave unit to become master, and click **Make Master**.

Executing a Command Cluster-Wide

To send a command to all members in the cluster, or to a specific member, perform the following steps. Sending a **show** command to all members collects all output and displays it on the console of the current unit. Other commands, such as **capture** and **copy**, can also take advantage of cluster-wide execution.

Detailed Steps

Step 1 Choose **Tools > Command Line Interface**.

Step 2 Enter the following command:

Command	Purpose
cluster exec [**unit** *unit_name*] *command* **Example:** `hostname# cluster exec show xlate`	Sends a command to all members, or if you specify the unit name, a specific member. To view member names, enter **cluster exec unit ?** (to see all names except the current unit), or enter the **show cluster info** command.

Step 3 Click **Send**.

Examples

To copy the same capture file from all units in the cluster at the same time to a TFTP server, enter the following command on the master unit:

```
cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each unit, are copied to the TFTP server. The destination capture file name is automatically attached with the unit name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster unit names.

The following sample output for the **cluster exec show port-channel** summary command shows EtherChannel information for each member in the cluster:

```
cluster exec show port-channel summary
primary(LOCAL):*********************************************************
 Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+----------+-----------------------------------------------
1         Po1             LACP      Yes  Gi0/0(P)
2         Po2             LACP      Yes  Gi0/1(P)
 secondary:*****************************************************************
 Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster   Ports
------+-------------+----------+-----------------------------------------------
1         Po1             LACP      Yes   Gi0/0(P)
2         Po2             LACP      Yes   Gi0/1(P)
```

Monitoring the ASA Cluster

- Cluster Dashboards, page 10-63
- Monitoring Screens, page 10-64
- Related Features, page 10-65

Cluster Dashboards

On the Home page on the master unit, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard. For more information, see the "Cluster Dashboard Tab" section on page 4-25 and the "Cluster Firewall Dashboard Tab" section on page 4-26.

Monitoring Screens

- Viewing the Cluster Summary, page 10-64
- Monitoring Cluster Resources, page 10-65
- Monitoring Cluster Traffic, page 10-65
- Monitoring the Cluster Control Link, page 10-65

Viewing the Cluster Summary

Choose **Monitoring > ASA Cluster > Cluster Summary**. This pane shows cluster information about the unit to which you are connected, as well as other units in the cluster. You can also change the master unit from this pane.

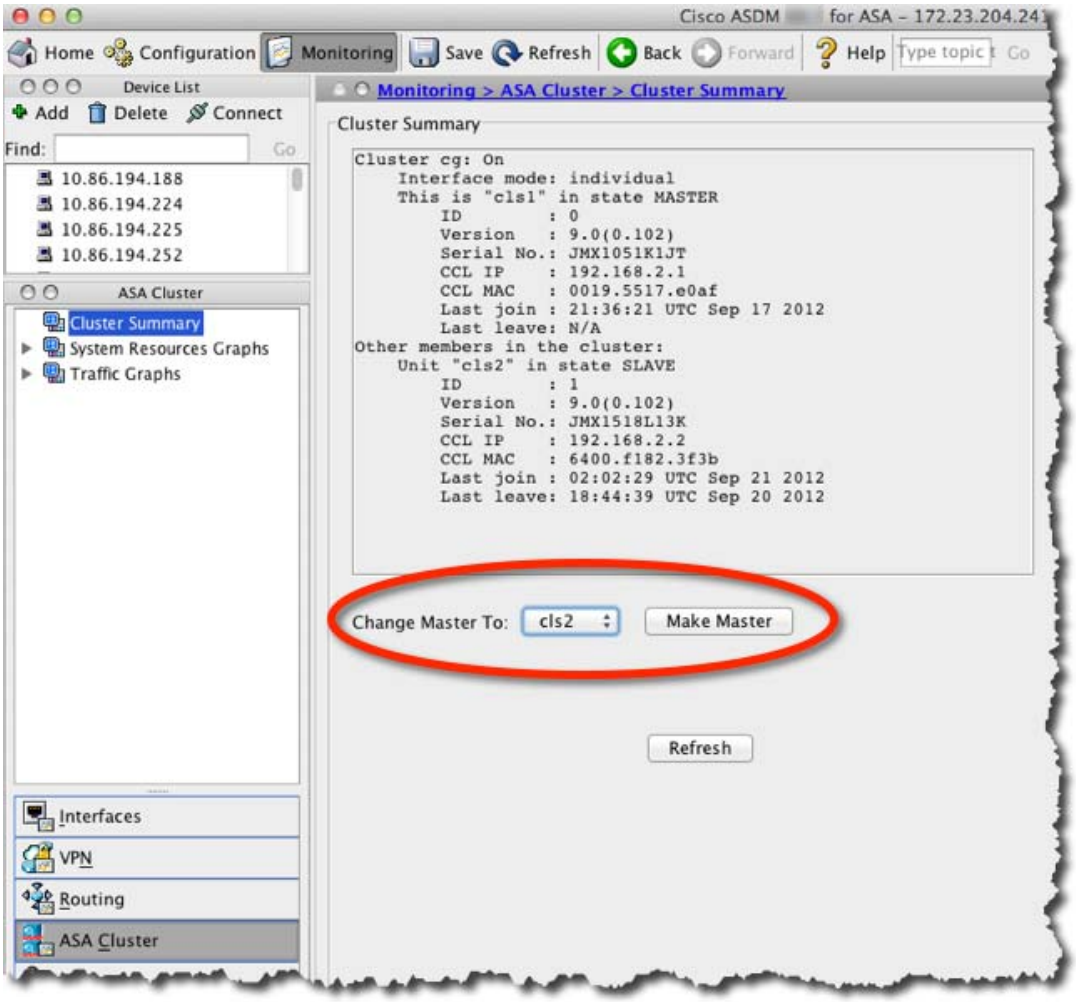

Monitoring Cluster Resources

CPU

Choose **Monitoring > ASA Cluster > System Resources Graphs > CPU**. This pane lets you create graphs or tables showing the CPU utilization across the cluster members.

Memory

Choose **Monitoring > ASA Cluster > System Resources Graphs > Memory**. This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster members.

Monitoring Cluster Traffic

Connections

Choose **Monitoring > ASA Cluster > Traffic Graphs > Connections**. This pane lets you create graphs or tables showing the Connections across the cluster members.

Throughput

Choose **Monitoring > ASA Cluster > Traffic Graphs > Throughput**. This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

Monitoring the Cluster Control Link

Choose **Monitoring > Properties > System Resources Graphs > Cluster Control Link**. This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Related Features

Command	Purpose
Wizards > Packet Capture Wizard	To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the master unit, which is then automatically enabled on all of the slave units in the cluster. See the "Configuring and Running Captures with the Packet Capture Wizard" section on page 102-1.
Configuration > Device Management > Interfaces > Edit Interface > Advanced	Creates a MAC address pool for an individual inteface.

Command	Purpose
Configuration > Device Management > Management Access > Command Line (CLI) > CLI Prompt	Sets the CLI prompt to include the cluster unit name. See the "Customizing a CLI Prompt" section on page 52-6.
Configuration > Device Management > Logging > Syslog Setup	Each unit in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster. See the "Including the Device ID in Non-EMBLEM Format Syslog Messages" section on page 96-21.

Configuration Examples for ASA Clustering

This section includes the following topics:

- Firewall on a Stick, page 10-66
- Traffic Segregation, page 10-67
- Redundant Interface (PBR or ECMP), page 10-67
- Spanned EtherChannel With Backup Links, page 10-68

Firewall on a Stick

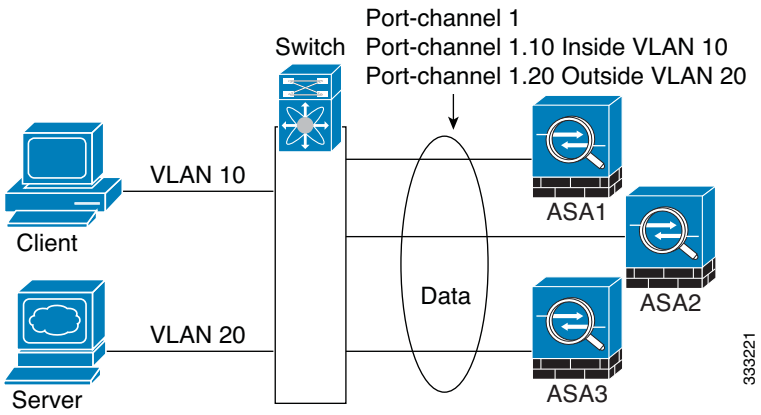

Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each ASA has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The ASA is the firewall between VLAN 10 and VLAN 20.

When using spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If an ASA becomes unavailable, the switch will rebalance traffic between the remaining units.

Traffic Segregation

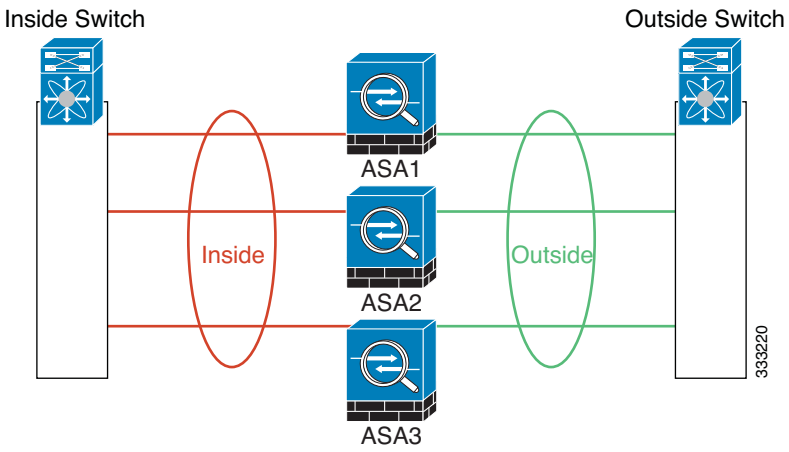

You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Redundant Interface (PBR or ECMP)

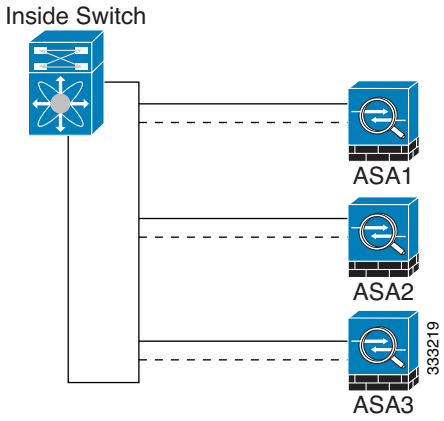

Redundant interfaces can be used to provide link-level redundancy.

When using individual interfaces, switching to a backup interface is similar to how it behaves in non-clustering mode. The ASA activates the backup link if the primary link fails. It takes time for the Spanning Tree on the switch to converge before the backup link is activated on the switch side. The backup links can be connected to a separate switch to provide inter-switch redundancy.

Spanned EtherChannel With Backup Links

The maximum number of active ports in an etherchannel is limited to 8 from the switch side. If you have an 8-ASA cluster, and you allocate 2 ports per unit to the EtherChannel, for a total of 16 ports total, then 8 of them have to be in standby mode. The ASA uses LACP to negotiate which links should be active or standby. If you enable multi-switch EtherChannel using VSS or vPC, you can achieve inter-switch redundancy. On the ASA, all physical ports are ordered first by the slot number then by the port number. In the following figure, the lower ordered port is the "primary" port (for example, GigabitEthernet 0/0), and the other one is the "secondary" port (for example, GigabitEthernet 0/1). You must guarantee symmetry in the hardware connection: all primary links must terminate on one switch, and all secondary links must terminate on another switch if VSS/vPC is used. The following diagram shows what happens when the total number of links grows as more units join the cluster:

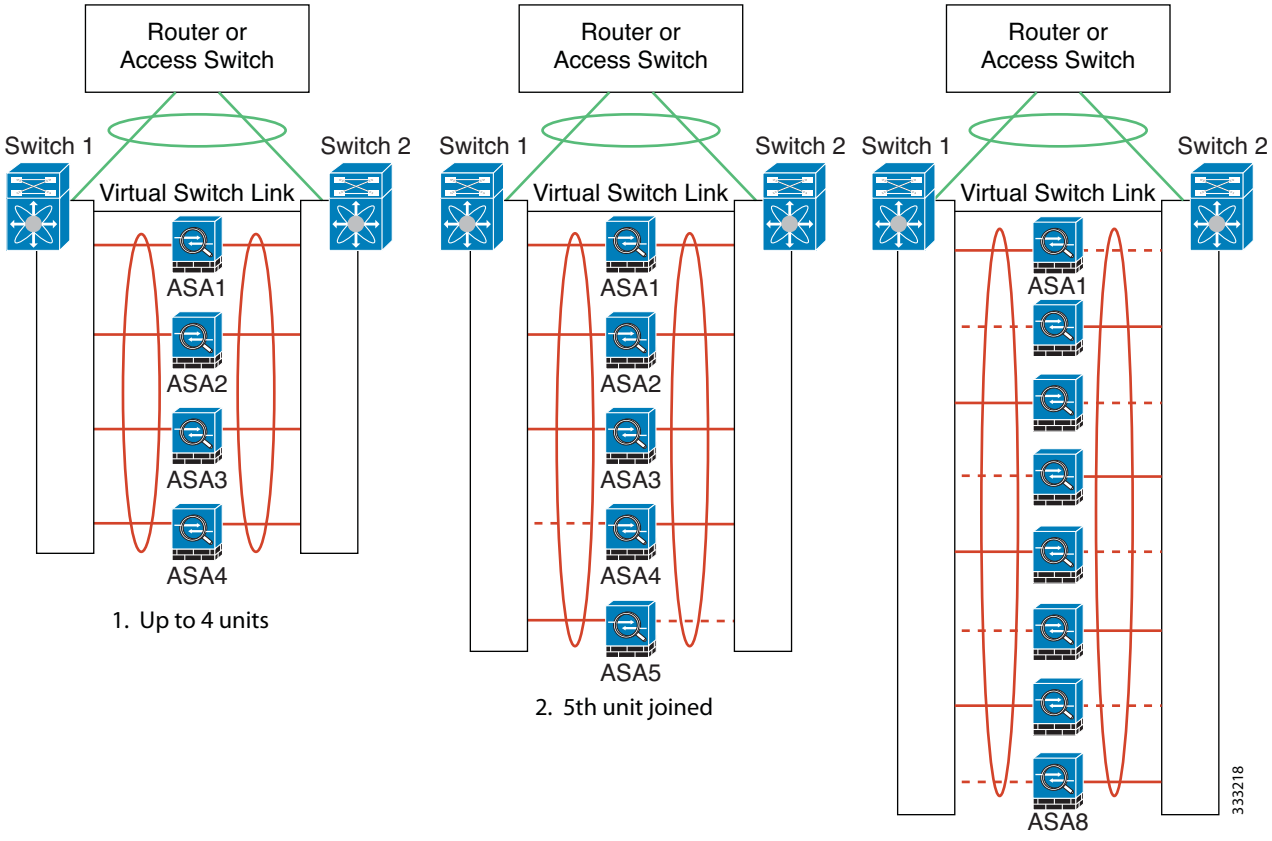

The principle is to first maximize the number of active ports in the channel, and secondly keep the number of active primary ports and the number of active secondary ports in balance. Note that when a 5th unit joins the cluster, traffic is not balanced evenly between all units.

Link or device failure is handled with the same principle. You may end up with a less-than-perfect load balancing situation. The following figure shows a 4-unit cluster with a single link failure on one of the units.

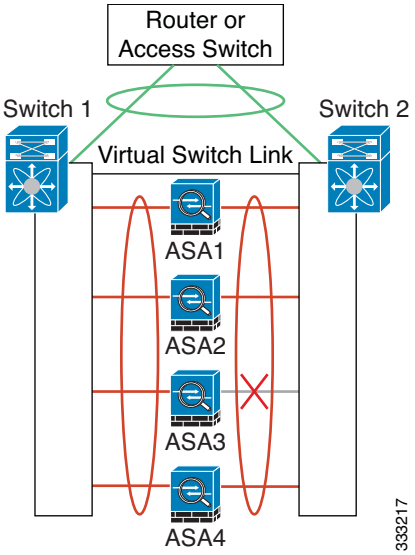

There could be multiple EtherChannels configured in the network. The following diagram shows an EtherChannel on the inside and one on the outside. An ASA is removed from the cluster if both primary and secondary links in one EtherChannel fail. This prevents the ASA from receiving traffic from the outside network when it has already lost connectivity to the inside network.

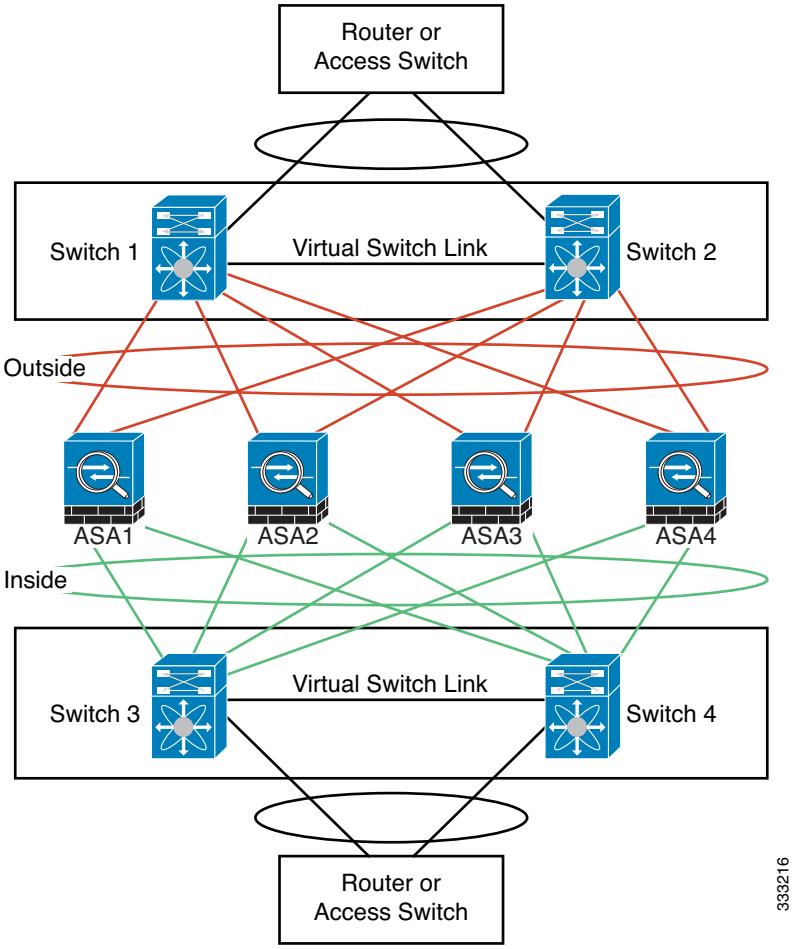

Feature History for ASA Clustering

Table 10-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 10-3 Feature History for Clustering

Feature Name	Platform Releases	Feature Information
ASA Clustering	9.0(1)	Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. We introduced or modified the following screens: Home > Device Dashboard Home > Cluster Dashboard Home > Cluster Firewall Dashboard Configuration > Device Management > Advanced > Address Pools > MAC Address Pools Configuration > Device Management > High Availability and Scalability > ASA Cluster Configuration > Device Management > Logging > Syslog Setup > Advanced Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced Configuration > Device Setup > Interfaces > Add/Edit Interface > IPv6 Configuration > Device Setup > Interfaces > Add/Edit EtherChannel Interface > Advanced Configuration > Firewall > Advanced > Per-Session NAT Rules Monitoring > ASA Cluster Monitoring > Properties > System Resources Graphs > Cluster Control Link Tools > Preferences > General Tools > System Reload Tools > Upgrade Software from Local Computer Wizards > High Availability and Scalability Wizard Wizards > Packet Capture Wizard Wizards > Startup Wizard

Information About Failover

This chapter provides an overview of the failover features that enable you to achieve high availability on the Cisco 5500 series ASAs. For information about configuring high availability, see Chapter 13, "Configuring Active/Active Failover" or Chapter 12, "Configuring Active/Standby Failover."

This chapter includes the following sections:

- Introduction to Failover and High Availability, page 11-1
- Failover System Requirements, page 11-2
- Failover and Stateful Failover Links, page 11-3
- Active/Active and Active/Standby Failover, page 11-8
- Stateless (Regular) and Stateful Failover, page 11-9
- Intra- and Inter-Chassis Module Placement for the ASA Services Module, page 11-11
- Transparent Firewall Mode Requirements, page 11-15
- Auto Update Server Support in Failover Configurations, page 11-16
- Failover Health Monitoring, page 11-18
- Failover Times, page 11-20
- Failover Messages, page 11-20

Introduction to Failover and High Availability

Configuring high availability requires two identical ASAs 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 if specific failover conditions are met. If those conditions are met, failover occurs.

The ASA supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover.

With Active/Active failover, both units can pass network traffic. This also lets you configure traffic sharing on your network. Active/Active failover is available only on units running in multiple context mode.

With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.

Both failover configurations support stateful or stateless (regular) failover.

Note When the ASA is configured for Active/Active Stateful Failover, you cannot enable IPsec or SSL VPN. Therefore, these features are unavailable. VPN failover is available for Active/Standby failover configurations only.

Failover System Requirements

This section describes the hardware, software, and license requirements for ASAs in a failover configuration.

This section includes the following topics:

- Hardware Requirements, page 11-2
- Software Requirements, page 11-2
- License Requirements, page 11-2

Hardware Requirements

The two units in a failover configuration must be the same model, have the same number and types of interfaces, the same SSMs installed (if any), and the same RAM installed.

If you are using units with different flash memory sizes in your failover configuration, make sure the unit with the smaller flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger flash memory to the unit with the smaller flash memory will fail.

Software Requirements

The two units in a failover configuration must be in the same operating modes (routed or transparent, single or multiple context). They must have the same major (first number) and minor (second number) software version. However, you can use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 8.3(1) to Version 8.3(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

License Requirements

The two units in a failover configuration do not need to have identical licenses; the licenses combine to make a failover cluster license. See the "Failover or ASA Cluster Licenses" section on page 5-31 for more information.

Failover and Stateful Failover Links

This section describes the failover and the Stateful Failover links, which are dedicated connections between the two units in a failover configuration. This section includes the following topics:

- Failover Link, page 11-3
- Stateful Failover Link, page 11-4
- Avoiding Interrupted Failover Links, page 11-5

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit. The following information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keep-alives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

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 ASA 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 ASA to terminate VPN tunnels.

You can use any unused interface on the device as the failover link; however, you cannot specify an interface that is currently configured with a name. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface should only be used for the failover link (and optionally for the Stateful Failover link).

Connect the failover link in one of the following two ways:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the failover interfaces of the ASA.
- Using a crossover Ethernet cable to connect the appliances directly, without the need for an external switch.

Note When you use a crossover cable for the failover link, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which interface failed and caused the link to come down.

Note The ASA supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

Although you can configure the failover and failover state links on a port channel link, this port channel cannot be shared with other firewall traffic.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link to pass all state information. You have three options for configuring a Stateful Failover link:

- You can use a dedicated Ethernet interface for the Stateful Failover link.
- You can share the failover link.
- You can share a regular data interface, such as the inside interface. However, this option is not recommended.

Connect a dedicated state link in one of the following two ways:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the failover interfaces of the ASA.
- Using a crossover Ethernet cable to connect the appliances directly, without the need for an external switch.

Note When you use a crossover cable for the state link, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which interface failed and caused the link to come down.

The ASA supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

Enable the PortFast option on Cisco switch ports that connect directly to the ASA.

If you use a data interface as the Stateful Failover link, you receive the following warning when you specify that interface as the Stateful Failover link:

```
******* WARNING ***** WARNING ******* WARNING ****** WARNING  *********
  Sharing Stateful failover interface with regular data interface is not
  a recommended configuration due to performance and security concerns.
******* WARNING ***** WARNING ******* WARNING ****** WARNING  *********
```

Sharing a data interface with the Stateful Failover interface can leave you vulnerable to replay attacks. Additionally, large amounts of Stateful Failover traffic may be sent on the interface, causing performance problems on that network segment.

Note Using a data interface as the Stateful Failover interface is supported in single context, routed mode only.

In multiple context mode, the Stateful Failover link resides in the system context. This interface and the failover interface are the only interfaces in the system context. All other interfaces are allocated to and configured from within security contexts.

Note The IP address and MAC address for the Stateful Failover link does not change at failover unless the Stateful Failover link is configured on a regular data interface.

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 ASA 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 ASA to terminate VPN tunnels.

Failover Interface Speed for Stateful Links

If you use the failover link as the Stateful Failover link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the Stateful Failover interface.

Use the following failover interface speed guidelines for the ASAs:

- Cisco ASA 5510
  - Stateful link speed can be 100 Mbps, even though the data interface can operate at 1 Gigabit due to the CPU speed limitation.
- Cisco ASA 5520/5540/5550
  - Stateful link speed should match the fastest data link.
- Cisco ASA 5580/5585-X
  - Use only non-management 1 Gigabit ports for the stateful link because management ports have lower performance and cannot meet the performance requirement for Stateful Failover.

For optimum performance when using long distance failover, the latency for the failover link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is more than10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

The ASA supports sharing of failover heartbeat and stateful link, but we recommend using a separate heartbeat link on systems with high Stateful Failover traffic.

Avoiding Interrupted Failover Links

Because the uses failover interfaces to transport messages between primary and secondary units, if a failover interface is down (that is, the physical link is down or the switch used to connect the interface is down), then the ASA failover operation is affected until the health of the failover interface is restored.

In the event that all communication is cut off between the units in a failover pair, both units go into the active state, which is expected behavior. When communication is restored and the two active units resume communication through the failover link or through any monitored interface, the primary unit remains active, and the secondary unit immediately returns to the standby state. This relationship is established regardless of the health of the primary unit.

Because of this behavior, stateful flows that were passed properly by the secondary active unit during the network split are now interrupted. To avoid this interruption, failover links and data interfaces should travel through different paths to decrease the chance that all links fail at the same time. In the event that only one failover link is down, the ASA takes a sample of the interface health, exchanges this information with its peer through the data interface, and performs a switchover if the active unit has a greater number of down interfaces. Subsequently, the failover operation is suspended until the health of the failover link is restored.

Depending upon their network topologies, several primary/secondary failure scenarios exist in ASA failover pairs, as shown in the following scenarios.

Scenario 1—Not Recommended

If a single switch or a set of switches are used to connect both failover and data interfaces between two ASAs, then when a switch or inter-switch-link is down, both ASAs become active. Therefore, the following two connection methods shown in Figure 11-1 and Figure 11-2 are NOT recommended.

Figure 11-1 Connecting with a Single Switch—Not Recommended

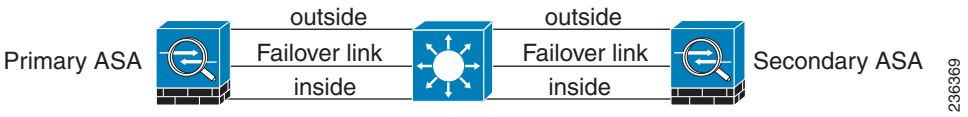

Figure 11-2 Connecting with a Double Switch—Not Recommended

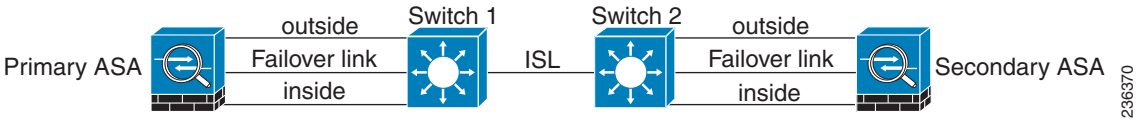

Scenario 2—Recommended

To make the ASA failover pair resistant to failover interface failure, we recommend that failover interfaces NOT use the same switch as the data interfaces, as shown in the preceding connections. Instead, use a different switch or use a direct cable to connect two ASA failover interfaces, as shown in Figure 11-3 and Figure 11-4.

Figure 11-3 Connecting with a Different Switch

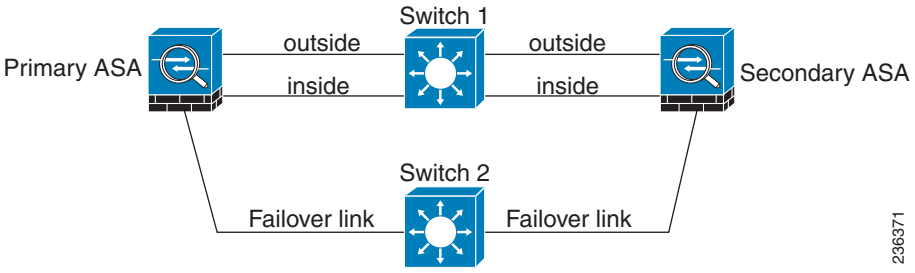

Figure 11-4 Connecting with a Cable

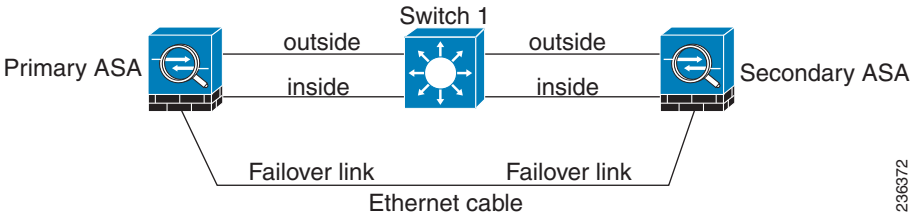

Scenario 3—Recommended

If the ASA data interfaces are connected to more than one set of switches, then a failover interface can be connected to one of the switches, preferably the switch on the secure side of network, as shown in Figure 11-5.

Figure 11-5 Connecting with a Secure Switch

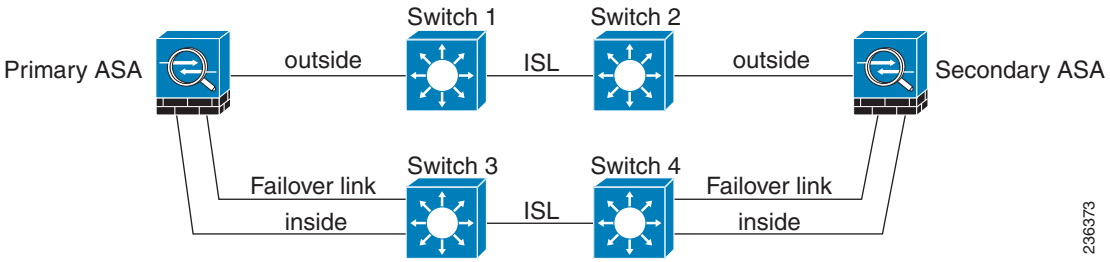

Scenario 4—Recommended

The most reliable failover configurations use a redundant interface on the failover interface, as shown in Figure 11-6 and Figure 11-7.

Figure 11-6 Connecting with Redundant Interfaces

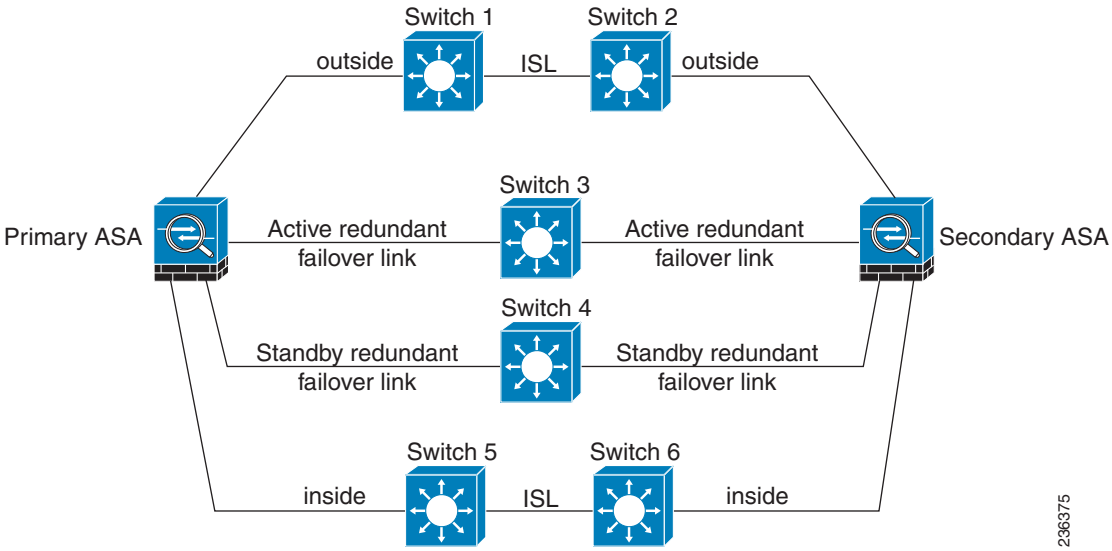

Figure 11-7 Connecting with Inter-switch Links

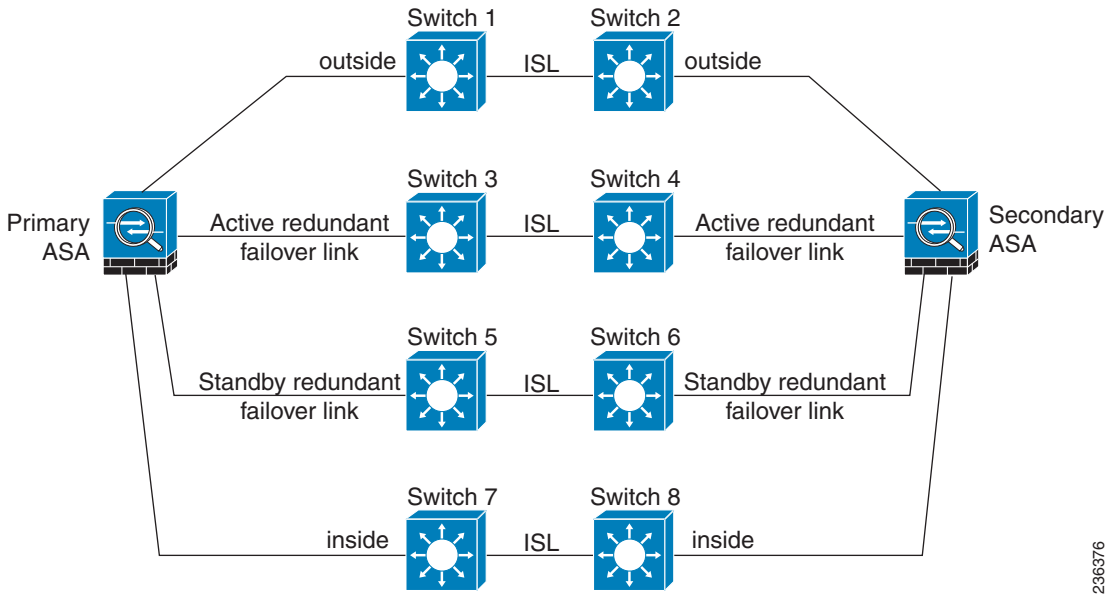

Active/Active and Active/Standby Failover

Two types of failover configurations are supported by the ASA: Active/Standby and Active/Active.

In Active/Standby failover, one unit is the active unit. It passes traffic. The standby unit does not actively pass traffic. When a failover occurs, the active unit fails over to the standby unit, which then becomes active. You can use Active/Standby failover for ASAs in single or multiple context mode, although it is most commonly used for ASAs in single context mode.

Active/Active failover is only available to ASAs in multiple context mode. In an Active/Active failover configuration, both ASAs can pass network traffic. In Active/Active failover, you divide the security contexts on the ASA into *failover groups*. A failover group is simply a logical group of one or more security contexts. Each group is assigned to be active on a specific ASA in the failover pair. When a failover occurs, it occurs at the failover group level.

For more detailed information about each type of failover, refer the following information:

- Chapter 12, “Configuring Active/Standby Failover”
- Chapter 13, “Configuring Active/Active Failover”

Determining Which Type of Failover to Use

The type of failover you choose depends upon your ASA configuration and how you plan to use the ASAs.

If you are running the ASA in single mode, then you can use only Active/Standby failover. Active/Active failover is only available to ASAs running in multiple context mode.

Note The ASA 5505 does not support multiple context mode or Active/Active failover.

VPN is not supported in multiple context mode or Active/Active failover.

If you are running the ASA in multiple context mode, then you can configure either Active/Active failover or Active/Standby failover.

- To allow both members of the failover pair to share the traffic, use Active/Active failover. Do not exceed 50% load on each device.
- If you do not want to share the traffic in this way, use Active/Standby or Active/Active failover.

Table 11-1 provides a comparison of some of the features supported by each type of failover configuration:

Table 11-1 Failover Configuration Feature Support

Feature	Active/Active	Active/Standby
Single Context Mode	No	Yes
Multiple Context Mode	Yes	Yes
Traffic Sharing Network Configurations	Yes	No
Unit Failover	Yes	Yes
Failover of Groups of Contexts	Yes	No
Failover of Individual Contexts	No	No

Stateless (Regular) and Stateful Failover

The ASA supports two types of failover, regular and stateful. This section includes the following topics:

- Stateless (Regular) Failover, page 11-9
- Stateful Failover, page 11-10

Stateless (Regular) Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Note Some configuration elements for clientless SSL VPN (such as bookmarks and customization) use the VPN failover subsystem, which is part of Stateful Failover. You must use Stateful Failover to synchronize these elements between the members of the failover pair. Stateless (regular) failover is not recommended for clientless SSL VPN.

Stateful Failover

When Stateful Failover is enabled, the active unit 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.

In Version 8.4 and later, Stateful Failover participates in dynamic routing protocols, like OSPF and EIGRP, so routes that are learned through dynamic routing protocols on the active unit are maintained in a Routing Information Base (RIB) table on the standby unit. Upon a failover event, packets travel normally with minimal disruption to traffic because the Active secondary ASA initially has rules that mirror the primary ASA. Immediately after failover, the re-convergence timer starts on the newly Active unit. Then the epoch number for the RIB table increments. During re-convergence, OSPF and EIGRP routes become updated with a new epoch number. Once the timer is expired, stale route entries (determined by the epoch number) are removed from the table. The RIB then contains the newest routing protocol forwarding information on the newly Active unit.

The following state information is passed to the standby ASA when Stateful Failover is enabled:

- NAT translation table
- TCP connection states
- UDP connection states
- The ARP table
- The Layer 2 bridge table (when running in transparent firewall mode)
- The HTTP connection states (if HTTP replication is enabled)
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signalling sessions
- ICMP connection state. ICMP connection replication is enabled only if the respective interface is assigned to an asymmetric routing group.

The following state information is *not* passed to the standby ASA when Stateful Failover is enabled:

- The HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- Inspected protocols are subject to advanced TCP-state tracking, and the TCP state of these connections is not automatically replicated. While these connections are replicated to the standby unit, there is a best-effort attempt to re-establish a TCP state.
- DHCP server address leases.
- State information for modules.
- Stateful Failover for phone proxy. When the active unit goes down, the call fails, media stops flowing, and the phone should unregister from the failed unit and reregister with the active unit. The call must be re-established.

The following clientless SSL VPN features are not supported with Stateful Failover:

- Smart Tunnels
- Port Forwarding
- Plugins
- Java Applets

- IPv6 clientless or Anyconnect sessions
- Citrix authentication (Citrix users must reauthenticate after failover)

Note If failover occurs during an active Cisco IP SoftPhone session, the call remains active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client loses connection with the Cisco CallManager. This occurs because there is no session information for the CTIQBE hangup message on the standby unit. When the IP SoftPhone client does not receive a response back from the Call Manager within a certain time period, it considers the CallManager unreachable and unregisters itself.

For VPN failover, VPN end-users should not have to reauthenticate or reconnect the VPN session in the event of a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.

Intra- and Inter-Chassis Module Placement for the ASA Services Module

You can place the primary and secondary ASASMs within the same switch or in two separate switches. The following sections describe each option:

- Intra-Chassis Failover, page 11-11
- Inter-Chassis Failover, page 11-12

Intra-Chassis Failover

If you install the secondary ASASM in the same switch as the primary ASASM, you protect against module-level failure. To protect against switch-level failure, as well as module-level failure, see the "Inter-Chassis Failover" section on page 11-12.

Even though both ASASMs are assigned the same VLANs, only the active module takes part in networking. The standby module does not pass any traffic.

Figure 11-8 shows a typical intra-switch configuration.

Figure 11-8 Intra-Switch Failover

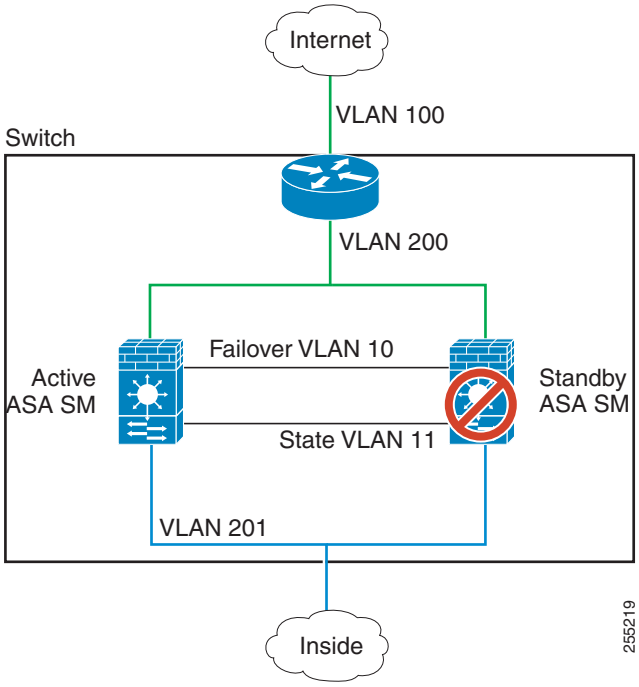

Inter-Chassis Failover

To protect against switch-level failure, you can install the secondary ASASM in a separate switch. The ASASM does not coordinate failover directly with the switch, but it works harmoniously with the switch failover operation. See the switch documentation to configure failover for the switch.

To accommodate the failover communications between ASASMs, we recommend that you configure a trunk port between the two switches that carries the failover and state VLANs. The trunk ensures that failover communication between the two units is subject to minimal failure risk.

For other VLANs, you must ensure that both switches have access to all firewall VLANs, and that monitored VLANs can successfully pass hello packets between both switches.

Figure 11-9 shows a typical switch and ASASM redundancy configuration. The trunk between the two switches carries the failover ASASM VLANs (VLANs 10 and 11).

Note ASASM failover is independent of the switch failover operation; however, ASASM works in any switch failover scenario.

Figure 11-9 Normal Operation

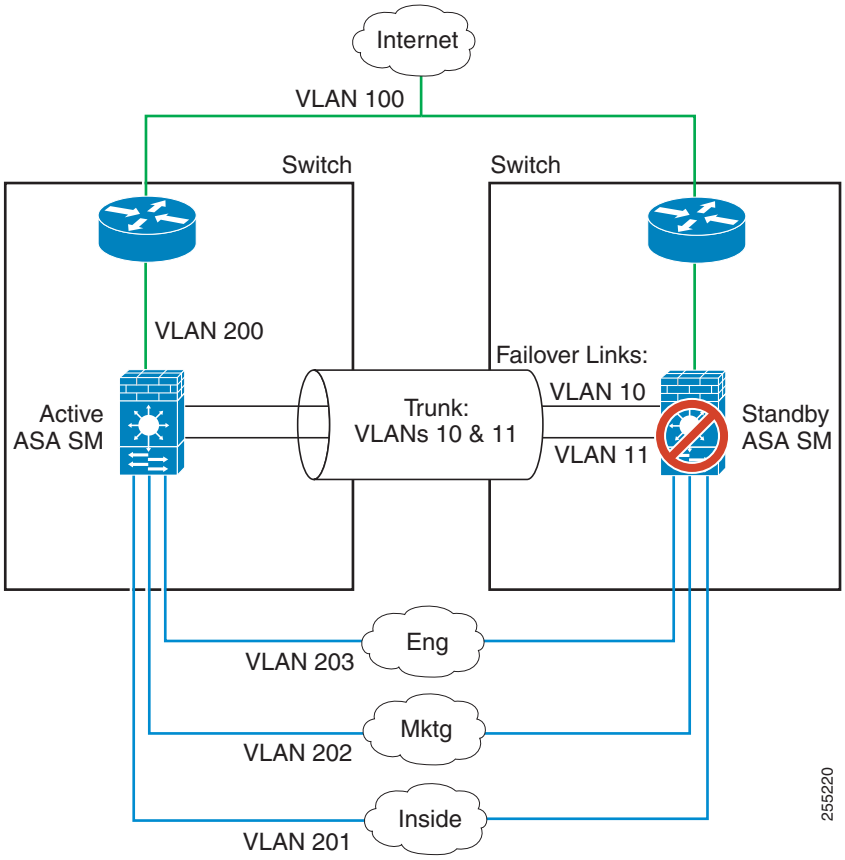

If the primary ASASM fails, then the secondary ASASM becomes active and successfully passes the firewall VLANs (Figure 11-10).

Figure 11-10 ASASM Failure

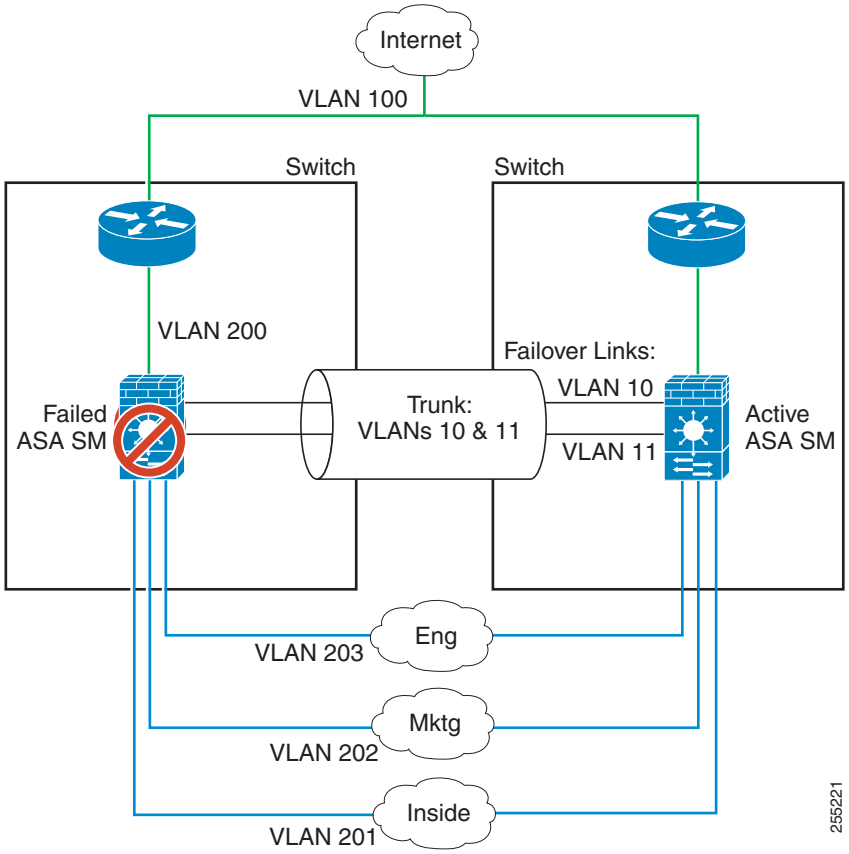

If the entire switch fails, as well as the ASASM (such as in a power failure), then both the switch and the ASASM fail over to their secondary units (Figure 11-11).

Figure 11-11 Switch Failure

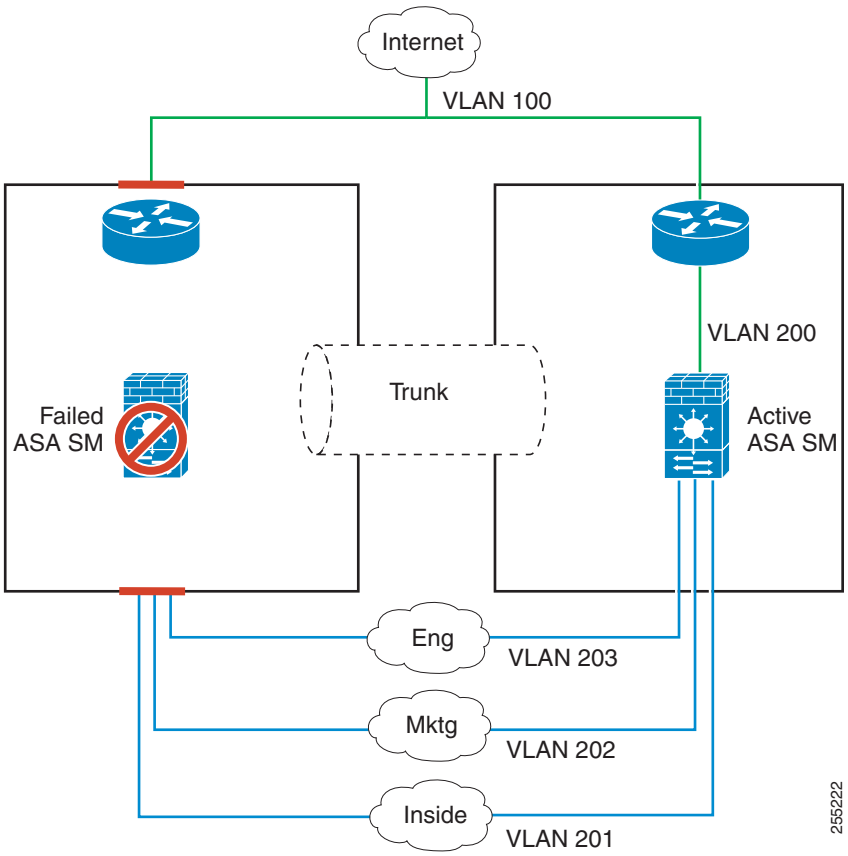

Transparent Firewall Mode Requirements

When the active unit fails over to the standby unit, the connected switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss while the port is in a blocking state, you can configure one of the following workarounds depending on the switch port mode:

- Access mode—Enable the STP PortFast feature on the switch:

  ```
  interface interface_id
  spanning-tree portfast
  ```

  The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.

- Trunk mode—Block BPDUs on the ASA on both the inside and outside interfaces:

  ```
  access-list id ethertype deny bpdu
  access-group id in interface inside_name
  access-group id in interface outside_name
  ```

Blocking BPDUs disables STP on the switch. Be sure not to have any loops involving the ASA in your network layout.

If neither of the above options are possible, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:

- Disable failover interface monitoring.
- Increase failover interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the failover interface holdtime.

Auto Update Server Support in Failover Configurations

You can use the Auto Update Server to deploy software images and configuration files to ASAs in an Active/Standby failover configuration. To enable Auto Update on an Active/Standby failover configuration, enter the Auto Update Server configuration on the primary unit in the failover pair.

The following restrictions and behaviors apply to Auto Update Server support in failover configurations:

- Only single mode, Active/Standby configurations are supported.
- When loading a new platform software image, the failover pair stops passing traffic.
- When using LAN-based failover, new configurations must not change the failover link configuration. If they do, communication between the units will fail.
- Only the primary unit will perform the call home to the Auto Update Server. The primary unit must be in the active state to call home. If it is not, the ASA automatically fails over to the primary unit.
- Only the primary unit downloads the software image or configuration file. The software image or configuration is then copied to the secondary unit.
- The interface MAC address and hardware-serial ID is from the primary unit.
- The configuration file stored on the Auto Update Server or HTTP server is for the primary unit only.

Auto Update Process Overview

The following is an overview of the Auto Update process in failover configurations. This process assumes that failover is enabled and operational. The Auto Update process cannot occur if the units are synchronizing configurations, if the standby unit is in the failed state for any reason other than SSM card failure, or if the failover link is down.

1. Both units exchange the platform and ASDM software checksum and version information.
2. The primary unit contacts the Auto Update Server. If the primary unit is not in the active state, the ASA first fails over to the primary unit and then contacts the Auto Update Server.
3. The Auto Update Server replies with software checksum and URL information.
4. If the primary unit determines that the platform image file needs to be updated for either the active or standby unit, the following occurs:
   a. The primary unit retrieves the appropriate files from the HTTP server using the URL from the Auto Update Server.
   b. The primary unit copies the image to the standby unit and then updates the image on itself.

c. If both units have new image, the secondary (standby) unit is reloaded first.

– If hitless upgrade can be performed when secondary unit boots, then the secondary unit becomes the active unit and the primary unit reloads. The primary unit becomes the active unit when it has finished loading.

– If hitless upgrade cannot be performed when the standby unit boots, then both units reload at the same time.

d. If only the secondary (standby) unit has new image, then only the secondary unit reloads. The primary unit waits until the secondary unit finishes reloading.

e. If only the primary (active) unit has new image, the secondary unit becomes the active unit, and the primary unit reloads.

f. The update process starts again at Step 1.

5. If the ASA determines that the ASDM file needs to be updated for either the primary or secondary unit, the following occurs:

   a. The primary unit retrieves the ASDM image file from the HTTP server using the URL provided by the Auto Update Server.

   b. The primary unit copies the ASDM image to the standby unit, if needed.

   c. The primary unit updates the ASDM image on itself.

   d. The update process starts again at Step 1.

6. If the primary unit determines that the configuration needs to be updated, the following occurs:

   a. The primary unit retrieves the configuration file from the using the specified URL.

   b. The new configuration replaces the old configuration on both units simultaneously.

   c. The update process begins again at Step 1.

7. If the checksums match for all image and configuration files, no updates are required. The process ends until the next poll time.

Monitoring the Auto Update Process

You can use the **debug auto-update client** or **debug fover cmd-exe** commands to display the actions performed during the Auto Update process. The following is sample output from the **debug auto-update client** command. Run **debug** commands from a terminal session.

```
Auto-update client: Sent DeviceDetails to /cgi-bin/dda.pl of server 192.168.0.21
Auto-update client: Processing UpdateInfo from server 192.168.0.21
   Component: asdm, URL: http://192.168.0.21/asdm.bint, checksum:
0x94bced0261cc992ae710faf8d244cf32
   Component: config, URL: http://192.168.0.21/config-rms.xml, checksum:
0x67358553572688a805a155af312f6898
   Component: image, URL: http://192.168.0.21/cdisk73.bin, checksum:
0x6d091b43ce96243e29a62f2330139419
Auto-update client: need to update img, act: yes, stby yes
name
ciscoasa(config)# Auto-update client: update img on stby unit...
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 2001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 2501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 3001, len = 1024
```

```
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 3501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 4001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 4501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 5001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 5501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 6001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 6501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 7001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 7501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 8001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 8501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 9001, len = 1024
auto-update: Fover file copy waiting at clock tick 6129280
fover_parse: Rcvd file copy ack, ret = 0, seq = 4
auto-update: Fover filecopy returns value: 0 at clock tick 6150260, upd time 145980 msecs
Auto-update client: update img on active unit...
fover_parse: Rcvd image info from mate
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
Beginning configuration replication: Sending to mate.
auto-update: HA safe reload: reload active waiting with mate state: 50
auto-update: HA safe reload: reload active waiting with mate state: 50

auto-update: HA safe reload: reload active waiting with mate state: 80
        Sauto-update: HA safe reload: reload active unit at clock tick: 6266860
Auto-update client: Succeeded: Image, version: 0x6d091b43ce96243e29a62f2330139419
```

The following syslog message is generated if the Auto Update process fails:

```
%ASA4-612002: Auto Update failed: file version: version reason: reason
```

The *file* is “image”, “asdm”, or “configuration”, depending on which update failed. The *version* is the version number of the update. And the *reason* is the reason that the update failed.

Failover Health Monitoring

The ASA monitors each unit for overall health and for interface health. See the following sections for more information about how the ASA performs tests to determine the state of each unit:

- Unit Health Monitoring, page 11-19
- Interface Monitoring, page 11-19

Unit Health Monitoring

The ASA determines the health of the other unit by monitoring the failover link. When a unit does not receive three consecutive hello messages on the failover link, the unit sends interface hello messages on each interface, including the failover interface, to validate whether or not the peer interface is responsive. The action that the ASA takes depends upon the response from the other unit. See the following possible actions:

- If the ASA receives a response on the failover interface, then it does not fail over.
- If the ASA does not receive a response on the failover link, but it does receive a response on another interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the ASA does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

You can configure the frequency of the hello messages and the hold time before failover occurs. A faster poll time and shorter hold time speed the detection of unit failures and make failover occur more quickly, but it can also cause "false" failures due to network congestion delaying the keepalive packets.

Interface Monitoring

You can monitor up to 250 interfaces divided between all contexts. You should monitor important interfaces. For example, you might configure one context to monitor a shared interface. (Because the interface is shared, all contexts benefit from the monitoring.)

When a unit does not receive hello messages on a monitored interface for half of the configured hold time, it runs the following tests:

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is operational, then the ASA performs network tests. The purpose of these tests is to generate network traffic to determine which (if either) unit has failed. At the start of each test, each unit clears its received packet count for its interfaces. At the conclusion of each test, each unit looks to see if it has received any traffic. If it has, the interface is considered operational. If one unit receives traffic for a test and the other unit does not, the unit that received no traffic is considered failed. If neither unit has received traffic, then the next test is used.
2. Network Activity test—A received network activity test. The unit counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the ARP test begins.
3. ARP test—A reading of the unit ARP cache for the 2 most recently acquired entries. One at a time, the unit sends ARP requests to these machines, attempting to stimulate network traffic. After each request, the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next machine. If at the end of the list no traffic has been received, the ping test begins.
4. Broadcast Ping test—A ping test that consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops.

If an interface has IPv4 and IPv6 addresses configured on it, the ASA uses the IPv4 addresses to perform the health monitoring.

If an interface has only IPv6 addresses configured on it, then the ASA uses IPv6 neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the ASA uses the IPv6 all nodes address (FE02::1).

If all network tests fail for an interface, but this interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed. If the threshold for failed interfaces is met, then a failover occurs. If the other unit interface also fails all the network tests, then both interfaces go into the "Unknown" state and do not count towards the failover limit.

An interface becomes operational again if it receives any traffic. A failed ASA returns to standby mode if the interface failure threshold is no longer met.

Note If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Failover Times

Table 11-2 shows the minimum, default, and maximum failover times.

Table 11-2 Cisco ASA 5500 Series ASA Failover Times

Failover Condition	Minimum	Default	Maximum
Active unit loses power or stops normal operation.	800 milliseconds	15 seconds	45 seconds
Active unit main board interface link down.	500 milliseconds	5 seconds	15 seconds
Active unit 4GE module interface link down.	2 seconds	5 seconds	15 seconds
Active unit IPS or CSC module fails.	2 seconds	2 seconds	2 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Failover Messages

When a failover occurs, both ASAs send out system messages. This section includes the following topics:

- Failover System Messages, page 11-20
- Debug Messages, page 11-21
- SNMP, page 11-21

Failover System Messages

The ASA issues a number of system messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the *syslog messages guide* guide. To enable logging, see Chapter 96, "Configuring Logging."

Note During switchover, failover logically shuts down and then bring up interfaces, generating syslog messages 411001 and 411002. This is normal activity.

Debug Messages

To see debug messages, enter the **debug fover** command. See the command reference for more information.

Note Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station. See Chapter 98, “Configuring SNMP” for more information.

Configuring Active/Standby Failover

This chapter describes how to configure Active/Standby failover and includes the following sections:

- Information About Active/Standby Failover, page 12-1
- Licensing Requirements for Active/Standby Failover, page 12-5
- Prerequisites for Active/Standby Failover, page 12-6
- Guidelines and Limitations, page 12-6
- Configuring Active/Standby Failover, page 12-7
- Controlling Failover, page 12-16
- Monitoring Active/Standby Failover, page 12-17
- Feature History for Active/Standby Failover, page 12-17

Information About Active/Standby Failover

This section describes Active/Standby failover and includes the following topics:

- Active/Standby Failover Overview, page 12-1
- Primary/Secondary Status and Active/Standby Status, page 12-2
- Device Initialization and Configuration Synchronization, page 12-2
- Command Replication, page 12-3
- Failover Triggers, page 12-4
- Failover Actions, page 12-4

Active/Standby Failover Overview

Active/Standby failover enables you to use a standby ASA to take over the functionality 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 and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note For multiple context mode, the ASA can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/Secondary Status and Active/Standby Status

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit is active and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both devices in the failover pair boot. Configurations are always synchronized from the active unit to the standby unit. When the standby unit completes its initial startup, it clears its running configuration (except for the failover commands needed to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit, and the secondary unit becomes the standby unit.

Note If the secondary unit boots without detecting the primary unit, it becomes the active unit. It uses its own MAC addresses for the active IP addresses. However, when the primary unit becomes available, the secondary unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. To avoid this, configure the failover pair with virtual MAC addresses. See the "Configuring Virtual MAC Addresses" section on page 12-15 for more information.

When the replication starts, the ASA console on the active unit displays the message "Beginning configuration replication: Sending to mate," and when it is complete, the ASA displays the message "End Configuration Replication to mate." During replication, commands entered on the active unit may not replicate properly to the standby unit, and commands entered on the standby unit may be overwritten by the configuration being replicated from the active unit. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

Note The **crypto ca server** command and related sub commands are not synchronized to the failover peer.

Note On the standby unit, the configuration exists only in running memory. To save the configuration to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**. Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

Command replication always flows from the active unit to the standby unit. As you apply your changes to the active unit in ASDM, the associated commands are sent across the failover link to the standby unit. You do not have to save the active configuration to flash memory to replicate the commands.

The following commands that are replicated to the standby ASA:

- All configuration commands except for **mode**, **firewall**, and **failover lan unit**
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands that are *not* replicated to the standby ASA:

- All forms of the **copy** command except for **copy running-config startup-config**
- All forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **show**
- **terminal pager** and **pager**

Note Changes made on the standby unit are not replicated to the active unit. If you enter a command on the standby unit, the ASA displays the message `**** WARNING **** Configuration Replication is NOT performed from Standby unit to Active unit. Configurations are no longer synchronized.` This message appears even when you enter many commands that do not affect the configuration.

Replicated commands are stored in the running configuration. To save replicated commands to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**.

Note Standby Failover does not replicate the following files and configuration components:

- AnyConnect images
- CSD images

- ASA images
- AnyConnect profiles
- Local Certificate Authorities (CAs)
- ASDM images

Failover Triggers

The unit can fail if one of the following events occurs:

- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- You force a failover. (See the "Forcing Failover" section on page 12-16.)

Failover Actions

In Active/Standby failover, failover occurs on a unit basis. Even on systems running in multiple context mode, you cannot fail over individual or groups of contexts.

Table 12-1 shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 12-1 Failover Behavior

Failure Event	Policy	Active Action	Standby Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover interface as failed	Mark failover interface as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Mark failover interface as failed	Become active	If the failover link is down at startup, both units become active.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.

Table 12-1 Failover Behavior (continued)

Failure Event	Policy	Active Action	Standby Action	Notes
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None.
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, then the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Optional Active/Standby Failover Settings

You can configure the following Active/Standby failover options when you initially configuring failover or after failover has been configured:

- HTTP replication with Stateful Failover—Allows connections to be included in the state information replication.
- Interface monitoring—Allows you to monitor up to 250 interfaces on a unit and control which interfaces affect your failover.
- Interface health monitoring—Enables the ASA to detect and respond to interface failures more quickly.
- Failover criteria setup—Allows you to specify a specific number of interfaces or a percentage of monitored interfaces that must fail before failover occurs.
- Virtual MAC address configuration—Ensures that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit.

Licensing Requirements for Active/Standby Failover

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASA 5505	Security Plus License. (Stateful failover is not supported).
ASA 5510, ASA 5512-X	Security Plus License.
All other models	Base License.

Failover units do not require the same license on each unit. If you have licenses on both units, they combine into a single running failover cluster license. The exceptions to this rule include:

- Security Plus license for the ASA 5505, 5510, and 5512-X—The Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.
- IPS module license for the ASA 5500-X—You must purchase an IPS module license for each unit, just as you would need to purchase a hardware module for each unit for other models.
- Encryption license—Both units must have the same encryption license.

Prerequisites for Active/Standby Failover

Active/Standby failover has the following prerequisites:

- Both units must be identical ASAs that are connected to each other through a dedicated failover link and, optionally, a Stateful Failover link.
- Both units must have the same software configuration and the proper license.
- Both units must be in the same mode (single or multiple, transparent or routed).

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, perform all steps in the system execution space unless otherwise noted.

Firewall Mode Guidelines

- Supported in transparent and routed firewall mode.

IPv6 Guidelines

- IPv6 failover is supported.

Model Guidelines

- Stateful failover is not supported on the ASA 5505.

Additional Guidelines and Limitations

Configuring port security on the switch(es) connected to an ASA failover pair can cause communication problems when a failover event occurs. This is because if a secure MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.

ASA failover replication fails if you try to make a configuration change in two or more contexts at the same time. The workaround is to make configuration changes on each unit sequentially.

The following guidelines and limitations apply for Active/Standby failover:

- To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.
- The standby IP addresses are used on the ASA that is currently the standby unit, and they must be in the same subnet as the active IP address on the corresponding interface on the active unit.
- If you change the console terminal pager settings on the active unit in a failover pair, the active console terminal pager settings change, but the standby unit settings do not. A default configuration issued on the active unit does affect behavior on the standby unit.
- When you enable interface monitoring, you can monitor up to 250 interfaces on a unit.
- By default, the ASA does not replicate HTTP session information when Stateful Failover is enabled. Because HTTP sessions are typically short-lived, and because HTTP clients typically retry failed connection attempts, not replicating HTTP sessions increases system performance without causing

serious data or connection loss. The **failover replication http** command enables the stateful replication of HTTP sessions in a Stateful Failover environment, but it could have a negative impact upon system performance.

- AnyConnect images must be the same on both ASAs in a failover pair. If the failover pair has mismatched images when a hitless upgrade is performed, then the WebVPN connection terminates in the final reboot step of the upgrade process, the database shows an orphaned session, and the IP pool shows that the IP address assigned to the client is “in use.”

Configuring Active/Standby Failover

This section describes how to configure Active/Standby failover. This section includes the following topics:

- Using the High Availability and Scalability Wizard, page 12-7
- Configuring Failover (Without the Wizard), page 12-10
- Configuring Optional Active/Standby Failover Settings, page 12-13

Using the High Availability and Scalability Wizard

The High Availability and Scalability Wizard guides you through a step-by-step process of creating either an Active/Active failover configuration, an Active/Standby failover configuration, or a VPN Cluster Load Balancing configuration.

As you go through the wizard, screens appear according to the type of failover that you are configuring and the hardware platform that you are using.

This section includes the following topics:

- Task Flow for Using the High Availability and Scalability Wizard, page 12-7
- Start the Wizard, page 12-8
- Failover Peer Connectivity and Compatibility Check, page 12-8
- Failover Link Configuration, page 12-8
- State Link Configuration, page 12-9
- Standby Address Configuration, page 12-10
- Summary, page 12-10

Task Flow for Using the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Standby failover using the High Availability and Scalability Wizard. Each step in the procedure corresponds to a wizard screen. Click **Next** after completing each step, except for the last one, before proceeding to the next step. Each step also includes a reference to additional information that you may need to complete the step.

Step 1 Start the wizard. See the “Start the Wizard” section on page 12-8.

Step 2 Enter the IP address of the failover peer on the Failover Peer Connectivity and Compatibility Check screen. Click **Test Compatibility**. You cannot move to the next screen until all compatibility tests have been passed. See the “Failover Peer Connectivity and Compatibility Check” section on page 12-8.

Step 3 Define the Failover Link in the Failover Link Configuration screen. See the "Failover Link Configuration" section on page 12-8.

Step 4 (Not available on the ASA 5505 ASA) Define the Stateful Failover link in the State Link Configuration screen. See the "State Link Configuration" section on page 12-9.

Step 5 Add standby addresses to the ASA interfaces in the Standby Address Configuration screen. See the "Standby Address Configuration" section on page 12-10.

Step 6 Review your configuration in the Summary screen. If necessary, click **Back** to go to a previous screen and make changes. See the "Summary" section on page 12-10.

Step 7 Click **Finish**. The failover configuration is sent to the ASA and to the failover peer.

Start the Wizard

Step 1 Choose **Wizards > High Availability and Scalability**.

Step 2 In the Configuration Type screen, click **Configure Active/Standby failover**.

Failover Peer Connectivity and Compatibility Check

The Failover Peer Connectivity and Compatibility Check 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.

To check failover peer connectivity and compatibility, perform the following steps:

Step 1 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. The field accepts both IPv4 and IPv6 addresses.

Step 2 Click **Next** 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 for the platform
- Software version compatibility
- Failover license compatibility
- Firewall mode compatibility (routed or transparent)
- Context mode compatibility (single or multiple)

Failover Link Configuration

The Failover Link Configuration screen appears *only* if you are configuring LAN-based failover.

To configure LAN-based failover, perform the following steps:

Step 1 Choose the LAN interface to use for failover communication from the drop-down list.

Step 2 Enter a name for the interface.

Step 3 Enter the IP address used for the failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 4 Enter the IP address used for the failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 5 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Step 6 (For 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 a default VLAN exists for every switch port, do not choose VLAN 1 for the inside port, because one less inside port will be available for another use.

Note To provide sufficient bandwidth for failover, do not use trunks or PoE for failover.

Step 7 (Optional) Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

State Link Configuration

Note The State Link Configuration screen does not appear on the ASA 5505.

The State Link Configuration screen lets you enable and disable Stateful Failover, and configure Stateful Failover link properties.

To enable Stateful Failover, perform the following steps:

Step 1 To pass state information across the LAN-based failover link, click **Use the LAN link as the State Link**.

Step 2 To disable Stateful Failover, click **Disable Stateful Failover**.

Step 3 To configure an unused interface as the Stateful Failover interface, click **Configure another interface for Stateful failover**.

Step 4 Choose the interface to use for Stateful Failover communication from the drop-down list.

Step 5 Enter the name for the Stateful Failover interface.

Step 6 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 7 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 8 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Standby Address Configuration

Use the Standby Address Configuration screen to assign standby IP addresses to the interface on the ASA. The interfaces currently configured on the failover devices appear. The interfaces are grouped by context, and the contexts are grouped by failover group.

To assign standby IP addresses to the interface on the ASA, perform the following steps:

Step 1 Click the plus sign (+) by a device name to display the interfaces on that device. Click the minus sign (-) by a device name to hide the interfaces on that device.

Step 2 Double-click the **Active IP** 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 failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 3 Double-click the **Standby IP** 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 failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 4 Check the **Is Monitored** check box to enable health monitoring for that interface. Uncheck the check box to disable health monitoring. By default, health monitoring of physical interfaces is enabled, and health monitoring of virtual interfaces is disabled.

Step 5 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.”

Summary

The Summary screen displays the results of the configuration steps that you performed in the previous wizard screens.

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** to return to the screen that you want to change. Make the change, and click **Next** until you return to the Summary screen.

Configuring Failover (Without the Wizard)

Follow these steps to configure Active/Standby failover on both units.

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.

Step 1 Choose the **Configuration > Device Management > Failover > Setup** tab.

Step 2 Check the **Enable Failover** check box.

Note Failover is not actually enabled until you apply your changes to the device.

Step 3 To encrypt the failover link, do the following:

a. (Optional) Check the **Use 32 hexadecimal character key** check box to enter a hexadecimal value for the encryption key in the Shared Key field.

b. Enter the encryption key in the Shared Key field.

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 the Use 32 hexadecimal character key check box is unchecked, 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.

Step 4 Select the interface to use for the failover link from the Interface list. 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.

Step 5 Specify the logical name of the interface used for failover communication in the Logical Name field.

Step 6 Specify the active IP address for the interface in the Active IP field. The IP address can be either an IPv4 or an IPv6 address. You cannot configure both types of addresses on the failover link interface.

Step 7 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 in the Subnet Mask/Prefix Length field. The name of the field changes depending upon the type of address specified in the Active IP field.

Step 8 Specify the IP address used by the secondary unit to communicate with the primary unit in the Standby IP field. The IP address can be an IPv4 or an IPv6 address.

Step 9 Select **Primary** or **Secondary** in the Preferred Role field to specify whether the preferred role for this ASA is as the primary or secondary unit.

Step 10 (Optional) Configure the Stateful Failover link by doing the following:

Note Stateful Failover is not available on the ASA 5505 platform. This area does not appear on ASDM running on an ASA 5505 ASA.

a. 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.

b. Specify the IP address for the Stateful Failover interface in the Active IP field. The IP address can be either an IPv4 or an IPv6 address. You cannot configure both types of addresses on the failover link interface. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.

c. Specify the mask (IPv4 address) or prefix (IPv6 address) for the Stateful Failover interface in the Subnet Mask/Prefix Length. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.

d. Specify the interface name used for failover communication in the Logical Name field. 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.

e. Specify the IP address used by the secondary unit to communicate with the primary unit in the Standby IP field. The IP address can be an IPv4 or an IPv6 address. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.

f. (Optional) Enable HTTP replication by checking the **Enable HTTP Replication** check box. This enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover.

Step 11 Click **Apply**.

The configuration is saved to the devices in the failover pair.

Configuring Interface Standby Addresses

Configuring standby IP address in ASDM changes depending upon the mode in which the unit is operating. This section includes the following topics:

- Configuring Interface Standby Addresses in Routed Firewall Mode, page 12-12
- Configuring the Management Interface Standby Address in Transparent Firewall Mode, page 12-13

Configuring Interface Standby Addresses in Routed Firewall Mode

To configure a standby address for each interface on the ASA, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. The IP address for each interface appears in the Active IP Address column. If configured, the standby IP address for the interface appears in the Standby IP address column. The failover interface and Stateful failover interface do not display IP address; you cannot change those address from this tab.

Step 2 For each interface that does not have a standby IP address, double-click the Standby IP Address field and do one of the following:

- Click the **...** button and select an IP address from the list.

- Type an IP address into the field. The address can be an IPv4 or an IPv6 address.

You can also specify whether or not the interface is monitored from this tab. For more information about configuring interface monitoring, see the "Disabling and Enabling Interface Monitoring" section on page 12-13.

Configuring the Management Interface Standby Address in Transparent Firewall Mode

If you are in multiple context mode, you must perform this procedure in each context.

To configure the management interface standby address on the ASA, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. Only the Management interface shows an IP address.

Step 2 For the Management interface that does not have a standby IP address, double-click the Standby IP Address field and do one of the following:

- Click the **...** button and select an IP address from the list.
- Type an IP address into the field. The address can be an IPv4 or an IPv6 address.

You can also specify whether or not the interface is monitored from this tab. For more information about configuring interface monitoring, see "Disabling and Enabling Interface Monitoring" section on page 12-13.

Configuring Optional Active/Standby Failover Settings

This section includes the following topics:

- Disabling and Enabling Interface Monitoring, page 12-13
- Configuring Failover Criteria, page 12-14
- Configuring the Unit and Interface Health Poll Times, page 12-15
- Configuring Virtual MAC Addresses, page 12-15

You can configure the optional Active/Standby failover settings when initially configuring the primary unit in a failover pair or on the active unit in the failover pair after the initial configuration.

Disabling and Enabling Interface Monitoring

You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This feature enables you to exclude interfaces attached to less critical networks from affecting your failover policy.

You can monitor up to 250 interfaces on a unit. By default, monitoring physical interfaces is enabled and monitoring subinterfaces is disabled.

Hello messages are exchanged during every interface poll frequency time period between the ASA failover pair. The failover interface poll time is 3 to 15 seconds. For example, if the poll time is set to 5 seconds, testing begins on an interface if 5 consecutive hellos are not heard on that interface (25 seconds).

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 or VLAN 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.

To enable or disable health monitoring for specific interfaces on units in single configuration mode, enter one of the following commands. Alternately, for units in multiple configuration mode, you must enter the commands within each security context.

To disable or enable monitoring of an interface, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. The Monitored column displays whether or not an interface is monitored as part of your failover criteria. If it is monitored, a check appears in the Monitored check box.

Step 2 To disable monitoring of a listed interface, uncheck the **Monitored** check box for the interface.

Step 3 To enable monitoring of a listed interface, check the **Monitored** check box for the interface.

Configuring Failover Criteria

You can specify a specific number of interface or a percentage of monitored interfaces that must fail before failover occurs. By default, a single interface failure causes failover.

Use the Configuration > Device Management > High Availability > Criteria 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.

For information about configuring the hold and poll times, see Configuring the Unit and Interface Health Poll Times, page 12-15.

To configure the interface policy, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Criteria** tab.

Step 2 In the Interface Policy area, do one of the following:

- To define a specific number of interfaces that must fail to trigger failover, enter a number from 1 to 250 in the Number of failed interfaces field. When the number of failed monitored interfaces exceeds the value you specify, the ASA fails over.
- To define a percentage of configured interfaces that must fail to trigger failover, enter a percentage in the Percentage of failed interfaces field. When the number of failed monitored interfaces exceeds the percentage you set, the ASA fails over.

Step 3 Click **Apply**.

Configuring the Unit and Interface Health Poll Times

The ASA sends hello packets out of each data interface to monitor interface health. The appliance sends hello messages across the failover link to monitor unit health. If the ASA does not receive a hello packet from the corresponding interface on the peer unit for over half of the hold time, then the additional interface testing begins. If a hello packet or a successful test result is not received within the specified hold time, the interface is marked as failed. Failover occurs if the number of failed interfaces meets the failover criteria.

Decreasing the poll and hold times enables the ASA to detect and respond to interface failures more quickly but may consume more system resources. Increasing the poll and hold times prevents the ASA from failing over on networks with higher latency.

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Criteria** tab.

Step 2 To configure the interface poll and hold times, change the following values in the Failover Poll Times area:

- 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.

Step 3 To configure the unit poll and hold times, change the following values in the Failover Poll Times area:

- 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.

Step 4 Click **Apply**.

Configuring Virtual MAC Addresses

The Configuration > Device Management > High Availability > MAC Addresses tab displays 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 the failover pair uses the burned-in NIC addresses as the MAC addresses.

Note You cannot configure a virtual MAC address for the failover or Stateful Failover links. The MAC and IP addresses for those links do not change during failover.

To configure the virtual MAC address for an interface, perform the following steps:

Step 1 Open the **Configuration > Device Management > High Availability > Failover > MAC Addresses** tab.

Step 2 To edit an existing virtual MAC address entry, double-click the row for the interface whose MAC addresses you want to change. To add a new virtual MAC address entry, click **Add**.

The Add/Edit Interface MAC Address dialog box appears.

Step 3 Type the new MAC address for the active interface in the Active MAC Address field.

Step 4 Type the new MAC address for the standby interface in the Standby MAC Address field.

Step 5 Click **OK**.

Step 6 To delete a virtual MAC address entry, perform the following steps:

a. Click the interface to select the table row.

b. Click **Delete**.

c. Click **OK**.

Controlling Failover

This sections describes how to control and monitor failover. This section includes the following topics:

- Forcing Failover, page 12-16
- Disabling Failover, page 12-17
- Restoring a Failed Unit, page 12-17

Forcing Failover

To force the standby unit to become active, perform the following steps:

Step 1 Choose **Monitoring > Properties > Failover > Status**.

Step 2 Click one of the following buttons:

- Click **Make Active** to make the unit the active unit.
- Click **Make Standby** to make the other unit in the pair the active unit.

Disabling Failover

To disable failover, perform the following steps:

Step 1 Choose **Configuration > Device Management > High Availability > Failover**.

Step 2 Uncheck the **Enable Failover** check box.

Restoring a Failed Unit

To restore a failed unit to an unfailed state, perform the following steps:

Step 1 Choose **Monitoring > Properties > Failover > Status**.

Step 2 Click **Reset Failover**.

Note

Monitoring Active/Standby Failover

Note After a failover event you should either re-launch ASDM or switch to another device in the Devices pane and then come back to the original ASA to continue monitoring the device. This action is necessary because the monitoring connection does not become re-established when ASDM is disconnected from and then reconnected to the device.

Choose **Monitoring > Properties > Failover** to monitor Active/Standby failover.

Feature History for Active/Standby Failover

Table 12-2 lists the release history for this feature.

Table 12-2 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
This feature was introduced.	7.0	This feature was introduced.
IPv6 support for failover added.	8.2(2)	We modified the following screens: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces

Configuring Active/Active Failover

This chapter describes how to configure Active/Active failover and includes the following sections:

- Information About Active/Active Failover, page 13-1
- Licensing Requirements for Active/Active Failover, page 13-6
- Prerequisites for Active/Active Failover, page 13-6
- Guidelines and Limitations, page 13-7
- Configuring Active/Active Failover, page 13-8
- Monitoring Active/Active Failover, page 13-22
- Feature History for Active/Active Failover, page 13-23

Information About Active/Active Failover

This section describes Active/Active failover. This section includes the following topics:

- Active/Active Failover Overview, page 13-1
- Primary/Secondary Status and Active/Standby Status, page 13-2
- Device Initialization and Configuration Synchronization, page 13-3
- Command Replication, page 13-3
- Failover Triggers, page 13-4
- Failover Actions, page 13-4

Active/Active Failover Overview

Active/Active failover is only available to ASAs in multiple context mode. In an Active/Active failover configuration, both ASAs can pass network traffic.

In Active/Active failover, you divide the security contexts on the ASA into *failover groups*. A failover group is simply a logical group of one or more security contexts. You can create a maximum of two failover groups. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

The failover group forms the base unit for failover in Active/Active failover. Interface failure monitoring, failover, and active/standby status are all attributes of a failover group rather than the unit. When an active failover group fails, it changes to the standby state while the standby failover group becomes

active. The interfaces in the failover group that becomes active assume the MAC and IP addresses of the interfaces in the failover group that failed. The interfaces in the failover group that is now in the standby state take over the standby MAC and IP addresses.

Note A failover group failing on a unit does not mean that the unit has failed. The unit may still have another failover group passing traffic on it.

When creating the failover groups, you should create them on the unit that will have failover group 1 in the active state.

Note Active/Active failover generates virtual MAC addresses for the interfaces in each failover group. 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.

Primary/Secondary Status and Active/Standby Status

As in Active/Standby failover, one unit in an Active/Active failover pair is designated the primary unit, and the other unit the secondary unit. Unlike Active/Standby failover, this designation does not indicate which unit becomes active when both units start simultaneously. Instead, the primary/secondary designation does two things:

- Determines which unit provides the running configuration to the pair when they boot simultaneously.
- Determines on which unit each failover group appears in the active state when the units boot simultaneously. Each failover group in the configuration is configured with a primary or secondary unit preference. You can configure 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, distributing the traffic across the devices.

  Note The ASA also provides load balancing, which is different from failover. Both failover and load balancing can exist on the same configuration. For information about load balancing, see the "Configuring Load Balancing" section on page 86-20.

Which unit each failover group becomes active on is determined as follows:

- When a unit boots while the peer unit is not available, both failover groups become active on the unit.
- When a unit boots while the peer unit is active (with both failover groups in the active state), the failover groups remain in the active state on the active unit regardless of the primary or secondary preference of the failover group until one of the following occurs:
  - A failover occurs.
  - You manually force a failover.

  - You configured preemption for the failover group, which causes the failover group to automatically become active on the preferred unit when the unit becomes available.
- When both units boot at the same time, each failover group becomes active on its preferred unit after the configurations have been synchronized.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both units in a failover pair boot. The configurations are synchronized as follows:

- When a unit boots while the peer unit is active (with both failover groups active on it), the booting unit contacts the active unit to obtain the running configuration regardless of the primary or secondary designation of the booting unit.
- When both units boot simultaneously, the secondary unit obtains the running configuration from the primary unit.

When the replication starts, the ASA console on the unit sending the configuration displays the message "Beginning configuration replication: Sending to mate," and when it is complete, the ASA displays the message "End Configuration Replication to mate." During replication, commands entered on the unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

Note On the unit receiving the configuration, the configuration exists only in running memory. To save the configuration to the flash memory on both units, select **File > Save Running Configuration to flash** from the menu bar in the system execution space on the unit that has failover group 1 in the active state. Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts configuration files from the disk on the primary unit to an external server, and then copy them to disk on the secondary unit, where they become available when the unit reloads.

Command Replication

After both units are running, commands are replicated from one unit to the other as follows:

- Changes 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.

- Changes 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.
- Changes 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 changes on the appropriate unit for command replication to occur causes the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

Table 13-1 lists the commands that are and are not replicated to the standby unit.

Table 13-1 Command Replication

Commands Replicated to the Standby Unit	Commands Not Replicated to the Standby Unit
All configuration commands except for **mode**, **firewall**, and **failover lan unit**	All forms of the **copy** command except for **copy running-config startup-config**
copy running-config startup-config	All forms of the **write** command except for **write memory**
delete	**debug**
mkdir	**failover lan unit**
rename	**firewall**
rmdir	**mode**
write memory	**show**

Failover Triggers

In Active/Active failover, failover can be triggered at the unit level if one of the following events occurs:

- The unit has a hardware failure.
- The unit has a power failure.
- The unit has a software failure.
- You force a failover. (See Forcing Failover, page 13-21.)

Failover is triggered at the failover group level when one of the following events occurs:

- Too many monitored interfaces in the group fail.
- You force a failover. (See Forcing Failover, page 13-21.)

You configure the failover threshold for each failover group by specifying the number or percentage of interfaces within the failover group that must fail before the group fails. Because a failover group can contain multiple contexts, and each context can contain multiple interfaces, it is possible for all interfaces in a single context to fail without causing the associated failover group to fail.

See the "Failover Health Monitoring" section on page 11-18 for more information about interface and unit monitoring.

Failover Actions

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, then 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.

Table 13-2 shows the failover action for each failure event. For each failure event, the policy (whether or not failover occurs), actions for the active failover group, and actions for the standby failover group are given.

Table 13-2 Failover Behavior for Active/Active Failover

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, the active failover group does not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless failover group preemption is configured, the failover groups remain active on their current unit.
Failover link failed at startup	No failover	Become active	Become active	If the failover link is down at startup, both failover groups on both units become active.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the failover interface as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Optional Active/Active Failover Settings

You can configure the following Active/Standby failover options when you initially configuring failover or after failover has been configured:

- Failover Group Preemption—Assigns a primary or secondary priority to a failover group to specify on which unit in the failover group becomes active when both units boot simultaneously.
- HTTP replication with Stateful Failover—Allows connections to be included in the state information replication.
- Interface monitoring—Allows you to monitor up to 250 interfaces on a unit and control which interfaces affect your failover.
- Interface health monitoring—Enables the security appliance to detect and respond to interface failures more quickly.
- Failover criteria setup—Allows you to specify a specific number of interfaces or a percentage of monitored interfaces that must fail before failover occurs.
- Virtual MAC address configuration—Ensures that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit.

Licensing Requirements for Active/Active Failover

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASA 5505	No support.
ASA 5510, ASA 5512-X	Security Plus License.
All other models	Base License.

Failover units do not require the same license on each unit. If you have licenses on both units, they combine into a single running failover cluster license. The exceptions to this rule include:

- Security Plus license for the ASA 5510 and 5512-X—The Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.
- IPS module license for the ASA 5500-X—You must purchase an IPS module license for each unit, just as you would need to purchase a hardware module for each unit for other models.
- Encryption license—Both units must have the same encryption license.

Prerequisites for Active/Active Failover

In Active/Active failover, both units must have the following:

- The same hardware model.
- The same number of interfaces.
- The same types of interfaces.

- The same software version, with the same major (first number) and minor (second number) version numbers. However you can use different versions of the software during an upgrade process; for example you can upgrade one unit from Version 7.0(1) to Version 7.9(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.
- The same software configuration.
- The same mode (multiple context mode).
- The proper license.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in multiple context mode only.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

IPv6 failover is supported.

Model Guidelines

Active/Active failover is not available on the Cisco ASA 5505.

Additional Guidelines and Limitations

No two interfaces in the same context should be configured in the same ASR group.

Configuring port security on the switch(es) connected to an ASA failover pair can cause communication problems when a failover event occurs. This is because if a secure MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.

ASA failover replication fails if you try to make a configuration change in two or more contexts at the same time. The workaround is to make configuration changes on each unit sequentially.

The following features are not supported for Active/Active failover:

- To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.
- The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.
- You can define a maximum number of two failover groups.
- Failover groups can only be added to the system context of devices that are configured for multiple context mode.
- You can create and remove failover groups only when failover is disabled.
- Entering the **failover group** command puts you in the failover group command mode. The **primary**, **secondary**, **preempt**, **replication http**, **interface-policy**, **mac address**, and **polltime interface** commands are available in the failover group configuration mode. Use the **exit** command to return to global configuration mode.

- The **failover polltime interface**, **failover interface-policy**, **failover replication http**, and **failover mac address** commands have no affect on Active/Active failover configurations. They are overridden by the following failover group configuration mode commands: **polltime interface**, **interface-policy**, **replication http**, and **mac address**.
- When removing failover groups, you must remove failover group 1 last. Failover group1 always contains the admin context. Any context not assigned to a failover group defaults to failover group 1. You cannot remove a failover group that has contexts explicitly assigned to it.
- VPN failover is unavailable. (It is available in Active/Standby failover configurations only.)

Configuring Active/Active Failover

This section describes how to configure Active/Active failover using an Ethernet failover link. When configuring LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

This section includes the following topics:

- Task Flow for Using the High Availability and Scalability Wizard, page 13-8
- Configuring Failover (Without the Wizard), page 13-12
- Configuring Asymmetric Routing Groups in Multiple Context Mode, page 13-20

Task Flow for Using the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Active failover using the High Availability and Scalability Wizard. Each step in the procedure corresponds to a wizard screen. Click **Next** after completing each step, except for the last one, before proceeding to the next step. Each step also includes a reference to additional information that you may need to complete the step.

Step 1 Start the wizard. See the "Start the Wizard" section on page 13-9.

Step 2 Enter the IP address of the failover peer in the Failover Peer Connectivity and Compatibility Check screen. Click **Test Compatibility**. You cannot move to the next screen until all compatibility tests have been passed. See the "Failover Peer Connectivity and Compatibility Check" section on page 13-9.

Step 3 If the ASA 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 ASA to multiple context mode, it reboots. ASDM automatically reestablishes communication with the ASA when it has finished rebooting. See the "Change a Device to Multiple Mode" section on page 13-9.

Step 4 Assign security contexts to failover groups in the Context Configuration screen. You can add and delete contexts in this screen. See the "Security Context Configuration" section on page 13-10.

Step 5 Define the Failover Link in the Failover Link Configuration screen. See the "Failover Link Configuration" section on page 13-10.

Step 6 Define the Stateful Failover link in the State Link Configuration screen. See the "State Link Configuration" section on page 13-10.

Step 7 Add standby addresses to the ASA interfaces in the Standby Address Configuration screen. See the "Standby Address Configuration" section on page 13-11.

Step 8 Review your configuration in the Summary screen. See "Summary" section on page 13-11.

Step 9 Click **Finish**. The failover configuration is sent to the ASA and to the failover peer.

Start the Wizard

Step 1 Choose **Wizards > High Availability and Scalability**.

Step 2 In the Configuration Type screen, click **Configure Active/Active failover**.

Failover Peer Connectivity and Compatibility Check

The Failover Peer Connectivity and Compatibility Check 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.

To check failover peer connectivity and compatibility, perform the following steps:

Step 1 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. The field accepts both IPv4 and IPv6 addresses.

Step 2 Click **Next** 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 for the platform
- Software version compatibility
- Failover license compatibility
- Firewall mode compatibility (routed or transparent)
- Context mode compatibility (single or multiple)

Change a Device to Multiple Mode

The Change Device to Multiple Mode dialog box requires that the ASA be in multiple context mode. This dialog box lets you convert a ASA in single context mode to multiple context mode.

When you convert from single context mode to multiple context mode, the ASA 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 ASA from single context mode to multiple context mode causes the ASA and its peer 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.

Note You must convert both the current ASA and its peer to multiple context mode before you can proceed.

To change the current ASA to multiple context mode, perform the following steps:

Step 1 Click **Change *device* To Multiple Context**, where *device* is the hostname of the ASA.

Step 2 Repeat this step for the peer ASA.

The status of the ASA appears during conversion to multiple context mode.

Security Context Configuration

The Security Context Configuration screen lets you assign security contexts to failover groups. It displays the name of currently configured security contexts, lets you add new ones, and change or remove existing ones as needed. In addition, it displays the failover group number to which the context is assigned and lets you change the failover group as needed. Although you can create security contexts in this screen, you cannot assign interfaces to those contexts or configure other properties for them. To configure context properties and assign interfaces to a context, choose **System > Security Contexts**.

Failover Link Configuration

The Failover Link Configuration screen appears *only* if you are configuring LAN-based failover.

To configure LAN-based failover, perform the following steps:

Step 1 Choose the LAN interface to use for failover communication from the drop-down list.

Step 2 Enter a name for the interface.

Step 3 Enter the IP address used for the failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 4 Enter the IP address used for the failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 5 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Step 6 (Optional) Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

State Link Configuration

The State Link Configuration screen lets you enable and disable Stateful Failover, and configure Stateful Failover link properties.

To enable Stateful Failover, perform the following steps:

Step 1 To pass state information across the LAN-based failover link, click **Use the LAN link as the State Link**.

Step 2 To disable Stateful Failover, click **Disable Stateful Failover**.

Step 3 To configure an unused interface as the Stateful Failover interface, click **Configure another interface for Stateful failover**.

Step 4 Choose the interface to use for Stateful Failover communication from the drop-down list.

Step 5 Enter the name for the Stateful Failover interface.

Step 6 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 7 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 8 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Standby Address Configuration

Use the Standby Address Configuration screen to assign standby IP addresses to the interface on the ASA. The interfaces currently configured on the failover devices appear. The interfaces are grouped by context, and the contexts are grouped by failover group.

To assign standby IP addresses to the interface on the ASA, perform the following steps:

Step 1 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.

Step 2 Double-click the **Active IP** 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 failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 3 Double-click the **Standby IP** 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 failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 4 Check the **Is Monitored** check box to enable health monitoring for that interface. Uncheck the check box to disable health monitoring. By default, health monitoring of physical interfaces is enabled, and health monitoring of virtual interfaces is disabled.

Step 5 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.”

Summary

The Summary screen displays the results of the configuration steps that you performed in the previous wizard screens.

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** to return to the screen that you want to change. Make the change, and click **Next** until you return to the Summary screen.

Configuring Failover (Without the Wizard)

- Failover-Multiple Mode, Security Context, page 13-12
- Failover - Routed, page 13-12
- Failover - Transparent, page 13-13
- Failover-Multiple Mode, System, page 13-14
- Failover > Setup Tab, page 13-14
- Failover > Criteria Tab, page 13-16
- Failover > Active/Active Tab, page 13-16
- Failover > MAC Addresses Tab, page 13-19

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 ASA 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.

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.

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.

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.

Failover-Multiple Mode, System

This pane includes tabs for configuring the system-level failover settings in the system context of an ASA 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 Chapter 11, "Information About 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 ASA 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 ASA, 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 ASA 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 ASA.

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 ASA 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.

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 ASA 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 ASA 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.

Failover > Active/Active Tab

Use this tab to enable Active/Active failover on the ASA by defining failover groups. In an Active/Active failover configuration, both ASAs pass network traffic. Active/Active failover is only available to ASAs in multiple mode.

A failover group is simply a logical group of security contexts. You can create two failover groups on the ASA. 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 ASA, 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 ASA 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 ASA.
  - 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.

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 ASA 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 ASA 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 ASA 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.

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).

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 ASA 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 ASA (usually primary).
  - Standby MAC Address—Identifies the MAC address on the standby ASA (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.

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 ASA (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 ASA (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Configuring Asymmetric Routing Groups in Multiple Context Mode

Note To configure asymmetric routing (ASR) groups, you must be in the admin context and it must be active.

To configure ASR groups, perform the following steps:

Step 1 Choose **Configuration > Device Management > Routing > ASR Groups**.

Step 2 Specify the ASR group IDs of the configured interfaces from the ASR Group ID drop-down list. The maximum number of groups that you can assign to one interface is eight. If other contexts have assigned interfaces, to a group, you may assign fewer groups to this context,

Step 3 Click **Apply** to save your changes to the running configuration.

Controlling Failover

This sections describes how to control and monitor failover. This section includes the following topics:

- Forcing Failover, page 13-21
- Disabling Failover, page 13-21
- Restoring a Failed Unit or Failover Group, page 13-22

Forcing Failover

To force failover at the unit level, follow these steps:

Step 1 Open **System > Monitoring > Failover > System**.

Step 2 Click one of the following buttons:

- Click **Make Active** to make the unit the active unit.
- Click **Make Standby** to make the other unit the active unit.

To force failover at the failover group level, follow these steps:

Step 1 Open **System > Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

Step 2 Cick one of the following buttons:

- Click **Make Active** to make the the failover group active on the security appliance.
- Click **Make Standby** to make the failover group active on the other security appliance.

Disabling Failover

Disabling failover on an Active/Active failover pair causes the failover groups to remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer. Enter the **no failover** command in the system execution space.

To disable failover, perform the following steps:

Step 1 Open the **System > Configuration> Device Setup > High Availability > Failover > Setup** tab.

Step 2 Clear the **Enable Failover** checkbox.

Restoring a Failed Unit or Failover Group

Restoring a failed unit or failover group moves the unit or failover group from the failed state to the standby state; it does not automatically make the failover group or unit active. Restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with failover preemption. If previously active, a failover group becomes active if it is configured with preemption and if the unit on which it failed is the preferred unit.

To restore a failed unit to an unfailed state, follow these steps:

Step 1 Open **System > Monitoring > Failover > System**.

Step 2 Click **Reset Failover**. Clicking this button on the active unit resets the standby unit.

To restore a failed failover group to an unfailed state, follow these steps:

Step 1 Open Open **System > Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to restore.

Step 2 Click **Reset Failover**.

Monitoring Active/Active Failover

Use the following screens in the Monitoring > Properties > Failover area to monitor Active/Active failover:

- System, page 13-22
- Failover Group 1 and Failover Group 2, page 13-23

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 ASA. The information shown is the same output you would receive from the **show failover** command. Refer to *Cisco ASA 5500 Series Command Reference* for more information about the displayed output.

The following actions are available on the System pane:

- Make Active—Click this button to make the ASA 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 ASA.

- Make Standby—Click this button to make the ASA 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 ASA.
- 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.

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.

You can perform the following actions from this pane:

- Make Active—Click this button to make the failover group active unit on the ASA.
- Make Standby—Click this button to force the failover group into the standby state on the ASA.
- 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.

Feature History for Active/Active Failover

Table 13-3 lists the release history for this feature.

Table 13-3 Feature History for Active/Active Failover

Feature Name	Releases	Feature Information
Active/Active failover	7.0	In an Active/Active failover configuration, both ASAs can pass network traffic. We introduced this feature and the relevant commands.
IPv6 Support in failover	8.2(2)	We modified the following screens: Configuration > Device Managment > High Availability > Failover > Setup Configuration > Device Managment > High Availability > Failover > Interfaces

PART 2

Configuring Interfaces

Starting Interface Configuration (ASA 5510 and Higher)

This chapter includes tasks for starting your interface configuration for the ASA 5510 and higher, including configuring Ethernet settings, redundant interfaces, and EtherChannels.

Note This chapter only applies to the ASA 5500 series appliances; for the ASASM, starting interface configuration consists of configuring switch ports and VLANs on the switch, and then assigning VLANs to the ASASM according to Chapter 2, "Configuring the Switch for Use with the ASA Services Module." To complete your interface configuration, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

For ASA 5505 configuration, see Chapter 16, "Starting Interface Configuration (ASA 5505)."

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 cluster interfaces, which have special requirements, see Chapter 10, "Configuring a Cluster of ASAs."

This chapter includes the following sections:

- Information About Starting ASA 5510 and Higher Interface Configuration, page 15-2
- Licensing Requirements for ASA 5510 and Higher Interfaces, page 15-8
- Guidelines and Limitations, page 15-10
- Default Settings, page 15-12
- Starting Interface Configuration (ASA 5510 and Higher), page 15-12
- Monitoring Interfaces, page 15-37
- Where to Go Next, page 15-41
- Feature History for ASA 5510 and Higher Interfaces, page 15-41

Information About Starting ASA 5510 and Higher Interface Configuration

This section includes the following topics:

- Auto-MDI/MDIX Feature, page 15-2
- Interfaces in Transparent Mode, page 15-2
- Management Interface, page 15-2
- Redundant Interfaces, page 15-5
- EtherChannels, page 15-5

Auto-MDI/MDIX Feature

For RJ-45 interfaces on the ASA 5500 series, 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.

Interfaces in Transparent Mode

Interfaces in transparent mode belong to a "bridge group," one bridge group for each network. You can have up to eight bridge groups of four interfaces each per context or in single mode. For more information about bridge groups, see the "Bridge Groups in Transparent Mode" section on page 18-2.

Management Interface

- Management Interface Overview, page 15-2
- Management Slot/Port Interface, page 15-3
- Using Any Interface for Management-Only Traffic, page 15-3
- Management Interface for Transparent Mode, page 15-4
- No Support for Redundant Management Interfaces, page 15-4
- Management 0/0 Interface on the ASA 5512-X through ASA 5555-X, page 15-4

Management Interface Overview

You can manage the ASA by connecting to:

- Any through-traffic interface
- A dedicated Management *Slot*/*Port* interface (if available for your model)

You may need to configure management access to the interface according to Chapter 52, "Configuring Management Access."

Management *Slot/Port* Interface

Table 15-1 shows the Management interfaces per model.-

Table 15-1 Management Interfaces Per Model

Model	Configurable for Through Traffic[1]	Management 0/0[2]	Management 0/1	Management 1/0	Management 1/1
ASA 5505	N/A	No	No	No	No
ASA 5510	**Yes**	**Yes**	No	No	No
ASA 5520	**Yes**	**Yes**	No	No	No
ASA 5540	**Yes**	**Yes**	No	No	No
ASA 5550	**Yes**	**Yes**	No	No	No
ASA 5580	**Yes**	**Yes**	**Yes**	No	No
ASA 5512-X	No	**Yes**	No	No	No
ASA 5515-X	No	**Yes**	No	No	No
ASA 5525-X	No	**Yes**	No	No	No
ASA 5545-X	No	**Yes**	No	No	No
ASA 5555-X	No	**Yes**	No	No	No
ASA 5585-X	**Yes**	**Yes**	**Yes**	**Yes**[3]	**Yes**[3]
ASASM	N/A	No	No	No	No

1. By default, the Management 0/0 interface is configured for management-only traffic. For supported models in routed mode, you can remove the limitation and pass through traffic. If your model includes additional Management interfaces, you can use them for through traffic as well. The Management interfaces might not be optimized for through-traffic, however.
2. The Management 0/0 interface is configured for ASDM access as part of the default factory configuration. See the "Factory Default Configurations" section on page 3-17 for more information.
3. If you installed an SSP in slot 1, then Management 1/0 and 1/1 provide management access to the SSP in slot 1 only.

Note If you installed an IPS module, then the IPS module management interface(s) provides management access for the IPS module only. For the ASA 5512-X through ASA 5555-X, the IPS SSP software module uses the same physical Management 0/0 interface as the ASA.

Using Any Interface for Management-Only Traffic

You can use any interface as a dedicated management-only interface by configuring it for management traffic, including an EtherChannel interface.

Management Interface for Transparent Mode

In transparent firewall mode, in addition to the maximum allowed through-traffic interfaces, you can also use the Management interface (either the physical interface, a subinterface (if supported for your model), or an EtherChannel interface comprised of Management interfaces (if you have multiple Management interfaces)) as a separate management interface. You cannot use any other interface types as management interfaces.

If your model does not include a Management interface, you must manage the transparent firewall from a data interface.

In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. To provide management per context, you can create subinterfaces of the Management interface and allocate a Management subinterface to each context. Note that the ASA 5512-X through ASA 5555-X do not allow subinterfaces on the Management interface, so for per-context management, you must connect to a data interface.

For 8.4(1) and later, the management interface is not part of a normal bridge group. Note that for operational purposes, it is part of a non-configurable bridge group.

Note In transparent firewall mode, the management interface updates the MAC address table in the same manner as a data interface; therefore you should not connect both a management and a data interface to the same switch unless you configure one of the switch ports as a routed port (by default Cisco Catalyst switches share a MAC address for all VLAN switch ports). Otherwise, if traffic arrives on the management interface from the physically-connected switch, then the ASA updates the MAC address table to use the *management* interface to access the switch, instead of the data interface. This action causes a temporary traffic interruption; the ASA will not re-update the MAC address table for packets from the switch to the data interface for at least 30 seconds for security reasons.

No Support for Redundant Management Interfaces

Redundant interfaces do not support Management *slot*/*port* interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.

Management 0/0 Interface on the ASA 5512-X through ASA 5555-X

The Management 0/0 interface on the ASA 5512-X through ASA 5555-X has the following characteristics:

- No through traffic support
- No subinterface support
- No priority queue support
- No multicast MAC support
- The IPS SSP software module shares the Management 0/0 interface. Separate MAC addresses and IP addresses are supported for the ASA and IPS module. You must perform configuration of the IPS IP address within the IPS operating system. However, physical characteristics (such as enabling the interface) are configured on the ASA.

Redundant Interfaces

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 ASA reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as device-level failover if desired.

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 the MAC Address and MTU" section on page 17-13 or the "Configuring Multiple Contexts" 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.

EtherChannels

An 802.3ad EtherChannel is a logical interface (called a port-channel interface) consisting of a bundle of individual Ethernet links (a channel group) so that you increase the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels.

This section includes the following topics:

- Channel Group Interfaces, page 15-5
- Connecting to an EtherChannel on Another Device, page 15-5
- Link Aggregation Control Protocol, page 15-6
- Load Balancing, page 15-7
- EtherChannel MAC Address, page 15-8

Channel Group Interfaces

Each channel group can have eight active interfaces. Note that you can assign up to 16 interfaces to a channel group. While only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The port is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and vlan numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the ASA EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Catalyst 6500 switch.

When the switch is part of a Virtual Switching System (VSS), then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch (see Figure 15-1).

Figure 15-1 Connecting to a VSS

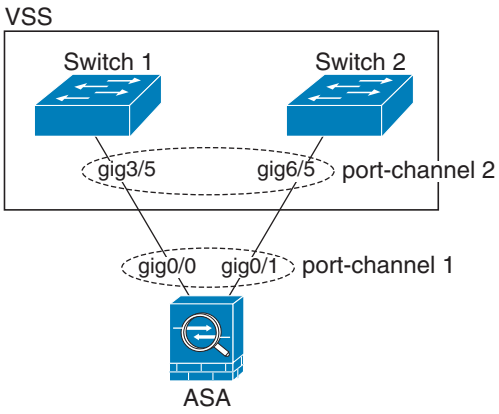

If you use the ASA in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS, one for each ASA (see Figure 15-1). On each ASA, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both ASAs (in this case, the EtherChannel will not be established because of the separate ASA system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby ASA.

Figure 15-2 Active/Standby Failover and VSS

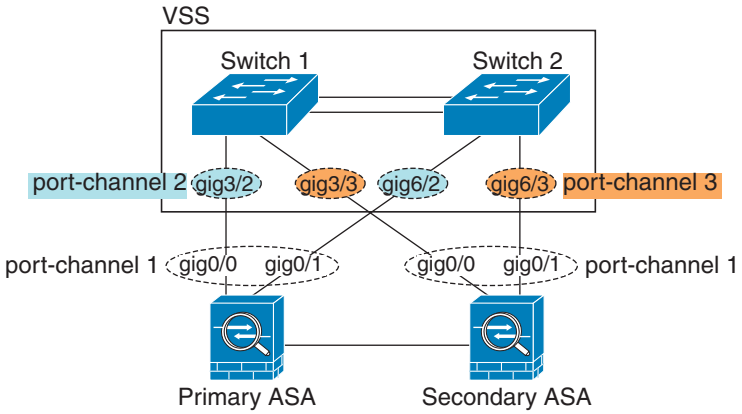

Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load Balancing

The ASA distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable; see the "Customizing the EtherChannel" section on page 15-31). The hash result is a 3-bit value (0 to 7).

The eight hash result values are distributed in a round robin fashion between the channel group interfaces, starting with the interface with the lowest ID (slot/port). For example, all packets with a hash result of 0 go to GigabitEthernet 0/0, packets with a hash result of 1 go to GigabitEthernet 0/1, packets with a hash result of 2 go to GigabitEthernet 0/2, and so on.

Because there are eight hash result values regardless of how many active interfaces are in the EtherChannel, packets might not be distributed evenly depending on the number of active interfaces.

Table 15-2 shows the load balancing amounts per interface for each number of active interfaces. The active interfaces in **bold** have even distribution.

Table 15-2 Load Distribution per Interface

# of Active Interfaces	% Distribution Per Interface							
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
1	**100%**	—	—	—	—	—	—	—
2	**50%**	**50%**	—	—	—	—	—	—
3	37.5%	37.5%	25%	—	—	—	—	—
4	**25%**	**25%**	**25%**	**25%**	—	—	—	—
5	25%	25%	25%	12.5%	12.5%	—	—	—
6	25%	25%	12.5%	12.5%	12.5%	12.5%	—	—
7	25%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	—
8	**12.5%**	**12.5%**	**12.5%**	**12.5%**	**12.5%**	**12.5%**	**12.5%**	**12.5%**

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC Address

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. In multiple context mode, you can automatically assign unique MAC addresses to interfaces, including an EtherChannel port interface. We recommend manually, or in multiple context mode, automatically configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

Licensing Requirements for ASA 5510 and Higher Interfaces

Model	License Requirement
ASA 5510	VLANs: Base License: 50 Security Plus License: 100 Interface Speed: Base License—All interfaces Fast Ethernet. Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet. Interfaces of all types[1]: Base License: 52 Security Plus License: 120
ASA 5520	VLANs: Base License: 150. Interfaces of all types[1]: Base License: 640
ASA 5540	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 840
ASA 5550	VLANs: Base License: 400 Interfaces of all types[1]: Base License: 1640

Model	License Requirement
ASA 5580	VLANs: Base License: 1024 Interfaces of all types[1]: Base License: 4176
ASA 5512-X	VLANs: Base License: 50 Interfaces of all types[1]: Base License: 328
ASA 5515-X	VLANs: Base License: 100 Interfaces of all types[1]: Base License: 528
ASA 5525-X	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 928
ASA 5545-X	VLANs: Base License: 300 Interfaces of all types[1]: Base License: 1328
ASA 5555-X	VLANs: Base License: 500 Interfaces of all types[1]: Base License: 2128
ASA 5585-X	VLANs: Base License: 1024 Interface Speed for SSP-10 and SSP-20: Base License—1-Gigabit Ethernet for fiber interfaces 10 GE I/O License (Security Plus)—10-Gigabit Ethernet for fiber interfaces (SSP-40 and SSP-60 support 10-Gigabit Ethernet by default.) Interfaces of all types[1]: Base License: 4176

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces.

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 15-12. Then, configure the logical interface parameters in the context execution space according to Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Firewall Mode Guidelines

- For transparent mode, you can configure up to eight bridge groups per context or for a single mode device.
- Each bridge group can include up to four interfaces.
- For multiple context, transparent mode, each context must use different interfaces; you cannot share an interface across contexts.

Failover Guidelines

- When you use a redundant or EtherChannel interface as a failover link, it must be pre-configured on both units in the failover pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the failover link itself is required for replication*.
- If you use a redundant or EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal.
- You can monitor redundant or EtherChannel interfaces for failover. When an active member interface fails over to a standby interface, this activity does not cause the redundant or EtherChannel interface to appear to be failed when being monitored for device-level failover. Only when all physical interfaces fail does the redundant or EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a failover or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link. To alter the configuration, you need to either shut down the EtherChannel while you make changes, or temporarily disable failover; either action prevents failover from occurring for the duration.
- Although you can configure the failover and failover state links on a port channel link, this port channel cannot be shared with other firewall traffic.

Clustering Guidelines

- When you use a redundant or EtherChannel interface as the cluster control link, it must be pre-configured on all units in the cluster; you cannot configure it on the primary unit and expect it to replicate to member units because *the cluster control link itself is required for replication*.
- To configure a spanned EtherChannel, see the "Configuring Spanned EtherChannels" section on page 10-37.
- To configure an individual cluster interface, see the "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32.

Redundant Interface Guidelines

- You can configure up to 8 redundant interface pairs.
- All ASA configuration refers to the logical redundant interface instead of the member physical interfaces.
- You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the ASA if they do not use the same physical interfaces.
- If you shut down the active interface, then the standby interface becomes active.
- Redundant interfaces do not support Management *slot*/*port* interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.
- For failover guidelines, see the "Failover Guidelines" section on page 15-10.
- For clustering guidelines, see the "Clustering Guidelines" section on page 15-10.

EtherChannel Guidelines

- You can configure up to 48 EtherChannels.
- Each channel group can have eight active interfaces. Note that you can assign up to 16 interfaces to a channel group. While only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.
- The device to which you connect the ASA EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Catalyst 6500 switch.
- The ASA does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the ASA will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch. In multiple context mode, these messages are not included in a packet capture, so you cannot diagnose the issue effectively.
- The ASA does not support connecting an EtherChannel to a switch stack. If the ASA EtherChannel is connected cross stack, and if the Master switch is powered down, then the EtherChannel connected to the remaining switch will not come up.
- All ASA configuration refers to the logical EtherChannel interface instead of the member physical interfaces.
- You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the ASA if they do not use the same physical interfaces.
- You cannot use interfaces on the 4GE SSM, including the integrated 4GE SSM in slot 1 on the ASA 5550, as part of an EtherChannel.
- For failover guidelines, see the "Failover Guidelines" section on page 15-10.
- For clustering guidelines, see the "Clustering Guidelines" section on page 15-10.

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 3-17.

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—Disabled.
- Redundant Interfaces—Enabled. However, for traffic to pass through the redundant interface, the member physical interfaces must also be enabled.
- Subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- EtherChannel port-channel interfaces—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces 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 (slot 1) 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 and 5585-X, the speed is set for automatic link negotiation.

Default Connector Type

The ASA 5550 (slot 1) and the 4GE SSM for the ASA 5510 and higher ASA include two connector types: copper RJ-45 and fiber SFP. RJ-45 is the default. You can configure the ASA 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 the following topics:

- Task Flow for Starting Interface Configuration, page 15-13
- Converting In-Use Interfaces to a Redundant or EtherChannel Interface, page 15-14
- Enabling the Physical Interface and Configuring Ethernet Parameters, page 15-23
- Configuring a Redundant Interface, page 15-25

- Configuring an EtherChannel, page 15-28
- Configuring VLAN Subinterfaces and 802.1Q Trunking, page 15-34
- Enabling Jumbo Frame Support (Supported Models), page 15-37

Task Flow for Starting Interface Configuration

Note If you have an existing configuration, and want to convert interfaces that are in use to a redundant or EtherChannel interface, perform your configuration offline using the CLI to minimize disruption. See the "Converting In-Use Interfaces to a Redundant or EtherChannel Interface" section on page 15-14.

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 15-23.

Physical interfaces are disabled by default.

Step 3 (Optional) Configure redundant interface pairs. See the "Configuring a Redundant Interface" section on page 15-25.

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 an EtherChannel. See the "Configuring an EtherChannel" section on page 15-28.

An EtherChannel groups multiple Ethernet interfaces into a single logical interface.

Note You cannot use interfaces on the 4GE SSM, including the integrated 4GE SSM in slot 1 on the ASA 5550, as part of an EtherChannel.

Step 5 (Optional) Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.

Step 6 (Optional) Enable jumbo frame support on the ASA 5580 and 5585-X according to the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37.

Step 7 (Multiple context mode only) To complete the configuration of interfaces in the system execution space, perform the following tasks that are documented in Chapter 9, "Configuring Multiple Context Mode":

- To assign interfaces to contexts, see the "Configuring a Security Context" section on page 9-20.
- (Optional) To automatically assign unique MAC addresses to context interfaces, see the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 9-24.

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 the MAC Address and MTU" section on page 17-13.

Step 8 Complete the interface configuration according to Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Converting In-Use Interfaces to a Redundant or EtherChannel Interface

If you have an existing configuration and want to take advantage of the redundant or EtherChannel interface feature for interfaces that are currently in use, you will have some amount of downtime when you convert to the logical interfaces.

This section provides an overview of how to convert your existing interfaces to a redundant or EtherChannel interface with minimal downtime. See the "Configuring a Redundant Interface" section on page 15-25 and the "Configuring an EtherChannel" section on page 15-28 fore more information.

- Detailed Steps (Single Mode), page 15-14
- Detailed Steps (Multiple Mode), page 15-19

Detailed Steps (Single Mode)

We recommend that you update your configuration offline as a text file, and reimport the whole configuration for the following reasons:

- Because you cannot add a named interface as a member of a redundant or EtherChannel interface, you must remove the name from the interface. When you remove the name from the interface, any command that referred to that name is deleted. Because commands that refer to interface names are widespread throughout the configuration and affect multiple features, removing a name from an in-use interface at the CLI or in ASDM would cause significant damage to your configuration, not to mention significant downtime while you reconfigure all your features around a new interface name.
- Changing your configuration offline lets you use the same interface names for your new logical interfaces, so you do not need to touch the feature configurations that refer to interface names. You only need to change the interface configuration.
- Clearing the running configuration and immediately applying a new configuration will minimize the downtime of your interfaces. You will not be waiting to configure the interfaces in real time.

Step 1 Connect to the ASA; if you are using failover, connect to the active ASA.

Step 2 If you are using failover, disable failover by choosing **Configuration > Device Management > High Availability > Failover** and unchecking the **Enable failover** check box. Click **Apply**, and continue at the warning.

Step 3 Copy the running configuration by choosing **Tools > Backup Configurations** and backing up the running configuration to your local computer. You can then expand the zip file and edit the running-config.cfg file with a text editor.

Be sure to save an extra copy of the old configuration in case you make an error when you edit it.

Step 4 For each in-use interface that you want to add to a redundant or EtherChannel interface, cut and paste all commands under the **interface** command to the end of the interface configuration section for use in creating your new logical interfaces. The only exceptions are the following commands, which should stay with the physical interface configuration:

- **media-type**

- **speed**
- **duplex**
- flowcontrol

Note You can only add *physical* interfaces to an EtherChannel or redundant interface; you cannot have VLANs configured for the physical interfaces.

Be sure to match the above values for all interfaces in a given EtherChannel or redundant interface. Note that the duplex setting for an EtherChannel interface must be Full or Auto.

For example, you have the following interface configuration. The bolded commands are the ones we want to use with three new EtherChannel interfaces, and that you should cut and paste to the end of the interface section.

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 no shutdown
!
interface Management0/1
 shutdown
 no nameif
```

```
 no security-level
 no ip address
```

Step 5 Above each pasted command section, create your new logical interfaces by entering one of the following commands:

- **interface redundant** *number* [1-8]
- **interface port-channel** *channel_id* [1-48]

For example:

```
...

interface port-channel 1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
 no shutdown
!
interface port-channel 2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface port-channel 3
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 no shutdown
```

Step 6 Assign the physical interfaces to the new logical interfaces:

- Redundant interface—Enter the following commands under the new **interface redundant** command:

  member-interface *physical_interface1*
  member-interface *physical_interface2*

  Where the physical interfaces are any two interfaces of the same type (either formerly in use or unused). You cannot assign a Management interface to a redundant interface.

  For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside redundant interfaces:

  ```
  interface redundant 1
   nameif outside
   security-level 0
   ip address 10.86.194.225 255.255.255.0
   member-interface GigabitEthernet0/0
   member-interface GigabitEthernet0/2

  interface redundant 2
   nameif inside
   security-level 100
   ip address 192.168.1.3 255.255.255.0
   member-interface GigabitEthernet0/1
   member-interface GigabitEthernet0/3
  ```

- EtherChannel interface—Enter the following command under each interface you want to add to the EtherChannel (either formerly in use or unused). You can assign up to 16 interfaces per EtherChannel, although only eight can be active; the others are in a standby state in case of failure.

  channel-group *channel_id* **mode active**

For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside EtherChannel interfaces:

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 shutdown
 no nameif
 no security-level
 no ip address

...
```

Step 7 Enable each formerly unused interface that is now part of a logical interface by adding **no** in front of the **shutdown** command.

For example, your final EtherChannel configuration is:

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 2 mode active
```

```
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface port-channel 1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
!
interface port-channel 2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
!
interface port-channel 3
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
```

Note Other optional EtherChannel parameters can be configured after you import the new configuration. See the "Configuring an EtherChannel" section on page 15-28.

Step 8 Save the entire new configuration, including the altered interface section.

Step 9 Re-zip the backup folder with the altered configuration.

Step 10 Choose **Tools > Restore Configurations**, and choose the altered configuration zip file. Be sure to replace the existing running configuration; do not merge them. See the "Restoring Configurations" section on page 101-25 for more information.

Step 11 Reenable failover by choosing **Configuration > Device Management > High Availability > Failover**, and checking the **Enable failover** check box. Click **Apply**, and click **No** when prompted if you want to configure basic failover settings.

Detailed Steps (Multiple Mode)

We recommend that you update your system and context configurations offline as text files, and reimport them for the following reasons:

- Because you cannot add an allocated interface as a member of a redundant or EtherChannel interface, you must deallocate the interface from any contexts. When you deallocate the interface, any context command that referred to that interface is deleted. Because commands that refer to interfaces are widespread throughout the configuration and affect multiple features, removing an allocation from an in-use interface at the CLI or in ASDM would cause significant damage to your configuration, not to mention significant downtime while you reconfigure all your features around a new interface.
- Changing your configuration offline lets you use the same interface names for your new logical interfaces, so you do not need to touch the feature configurations that refer to interface names. You only need to change the interface configuration.
- Clearing the running system configuration and immediately applying a new configuration will minimize the downtime of your interfaces. You will not be waiting to configure the interfaces in real time.

Step 1 Connect to the ASA, and change to the system; if you are using failover, connect to the active ASA.

Step 2 If you are using failover, disable failover by choosing **Configuration > Device Management > High Availability > Failover** and unchecking the **Enable failover** check box. Click **Apply**, and continue at the warning.

Step 3 In the system, copy the running configuration by choosing **File > Show Running Configuration in New Window** and copying the display output to a text editor.

Be sure to save an extra copy of the old configuration in case you make an error when you edit it.

For example, you have the following interface configuration and allocation in the system configuration, with shared interfaces between two contexts.

System

```
interface GigabitEthernet0/0
    no shutdown
interface GigabitEthernet0/1
    no shutdown
interface GigabitEthernet0/2
    shutdown
interface GigabitEthernet0/3
    shutdown
interface GigabitEthernet0/4
    shutdown
interface GigabitEthernet0/5
    shutdown
```

```
interface Management0/0
    no shutdown
interface Management1/0
    shutdown
!
context customerA
    allocate-interface gigabitethernet0/0 int1
    allocate-interface gigabitethernet0/1 int2
    allocate-interface management0/0 mgmt
context customerB
    allocate-interface gigabitethernet0/0
    allocate-interface gigabitethernet0/1
    allocate-interface management0/0
```

Step 4 Get copies of *all* context configurations that will use the new EtherChannel or redundant interface.See the "Backing Up and Restoring Configurations or Other Files" section on page 101-20.

For example, you download the following context configurations (interface configuration shown):

CustomerA Context

```
interface int1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
!
interface int2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface mgmt
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 management-only
```

CustomerB Context

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.20.15.5 255.255.255.0
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.6.78 255.255.255.0
!
interface Management0/0
 nameif mgmt
 security-level 100
 ip address 10.8.1.8 255.255.255.0
 management-only
```

Step 5 In the system configuration, create the new logical interfaces according to the "Configuring a Redundant Interface" section on page 15-25 or the "Configuring an EtherChannel" section on page 15-28. Be sure to enter the **no shutdown** command on any additional physical interfaces you want to use as part of the logical interface.

Note You can only add *physical* interfaces to an EtherChannel or redundant interface; you cannot have VLANs configured for the physical interfaces.

Be sure to match physical interface parameters such as speed and duplex for all interfaces in a given EtherChannel or redundant interface. Note that the duplex setting for an EtherChannel interface must be Full or Auto.

For example, the new configuration is:

System

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 no shutdown
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 no shutdown
!
interface port-channel 1
interface port-channel 2
interface port-channel 3
```

Step 6 Change the interface allocation per context to use the new EtherChannel or redundant interfaces. See the "Configuring a Security Context" section on page 9-20.

For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside redundant interfaces:

```
context customerA
    allocate-interface port-channel1 int1
    allocate-interface port-channel2 int2
    allocate-interface port-channel3 mgmt
context customerB
    allocate-interface port-channel1
    allocate-interface port-channel2
```

```
allocate-interface port-channel3
```

Note You might want to take this opportunity to assign mapped names to interfaces if you have not done so already. For example, the configuration for customerA does not need to be altered at all; it just needs to be reapplied on the ASA. The customerB configuration, however, needs to have all of the interface IDs changed; if you assign mapped names for customerB, you still have to change the interface IDs in the context configuration, but mapped names might help future interface changes.

Step 7 For contexts that do not use mapped names, change the context configuration to use the new EtherChannel or redundant interface ID. (Contexts that use mapped interface names do not require any alteration.)

For example:

CustomerB Context

```
interface port-channel1
 nameif outside
 security-level 0
 ip address 10.20.15.5 255.255.255.0
!
interface port-channel2
 nameif inside
 security-level 100
 ip address 192.168.6.78 255.255.255.0
!
interface port-channel3
 nameif mgmt
 security-level 100
 ip address 10.8.1.8 255.255.255.0
 management-only
```

Step 8 Copy the new context configuration files over the old ones. For example, for contexts in flash memory, in the system choose **Tools > File Management**, then choose **File Transfer > Between Local PC and Flash**. This tool lets you choose each configuration file and copy it to your local computer. This change only affects the startup configuration; the running configuration is still using the old context configuration.

Step 9 Copy the entire new system configuration to the clipboard, including the altered interface section.

Step 10 In ASDM, choose **Tools > Command Line Interface**, and click the **Multiple Line** radio button.

Step 11 Enter **clear configure all** as the first line, paste the new configuration after it, and click **Send**. The **clear** command clears the running configuration (both system and contexts), before applying the new configuration.

Traffic through the ASA stops at this point. All of the new context configurations now reload. When they are finished reloading, traffic through the ASA resumes.

Step 12 Close the Command Line Interface dialog box, and choose **File > Refresh ASDM with the Running Configuration**.

Step 13 Reenable failover by choosing **Configuration > Device Management > High Availability > Failover**, and checking the **Enable failover** check box. Click **Apply**, and click **No** when prompted if you want to configure basic failover settings.

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

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.

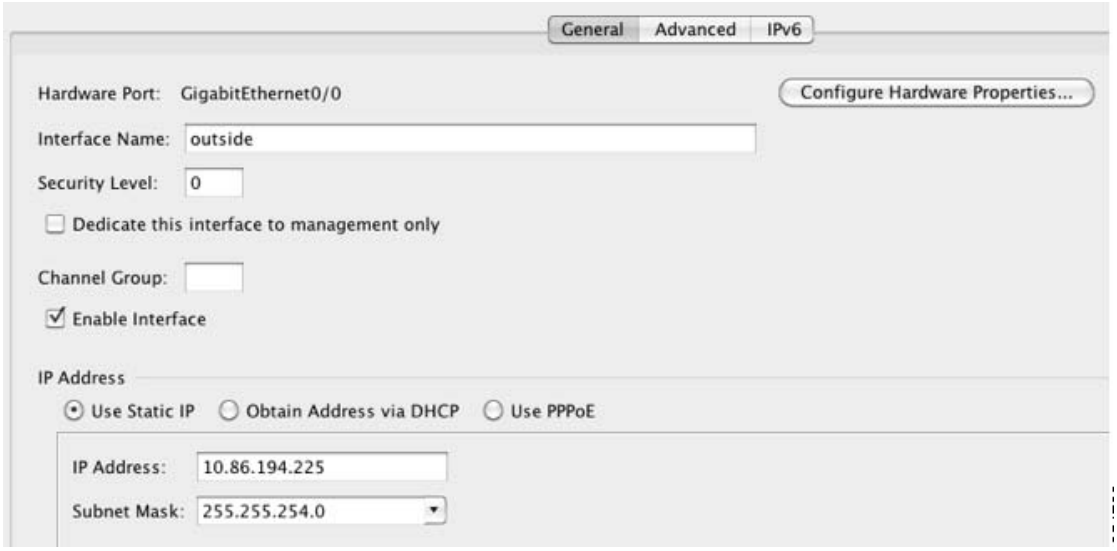

Note In single mode, this procedure only covers a subset of the parameters on the Edit Interface dialog box; to configure other parameters, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)." Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See the "Configuring Multiple Contexts" section on page 9-15.

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, speed, and enable pause frames for flow control, click **Configure Hardware Properties**.

a. For slot 1 on the ASA 5550 ASA or the 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.

Note The duplex setting for an EtherChannel interface must be Full or Auto.

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 ASA, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. See the "Auto-MDI/MDIX Feature" section on page 15-2.

d. To enable pause (XOFF) frames for flow control on 1-Gigabit and 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-water mark. The default *high_water* value is 128 KB (10 GigabitEthernet) and 24 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and 47 KB (1 GigabitEthernet). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark. By default, the *low_water* value is 64 KB (10 GigabitEthernet) and 16 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and

47 KB (1 GigabitEthernet). The link partner can resume traffic after receiving an XON, or after the XOFF expires, as controlled by the timer value in the pause frame. The default *pause_time* value is 26624; you can set it between 0 and 65535. If the buffer usage is consistently above the high-water mark, 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 15-25.
- Configure an EtherChannel. See the "Configuring an EtherChannel" section on page 15-28.
- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Configuring Multiple Contexts" section on page 9-15.
- For single context mode, complete the interface configuration. See Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

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 ASA 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 15-25
- Changing the Active Interface, page 15-28

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.
- Redundant interface delay values are configurable, but by default the ASA inherits the default delay values based on the physical type of its member interfaces.
- See also the "Redundant Interface Guidelines" section on page 15-11.

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

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 Choose **Add > Redundant Interface**.

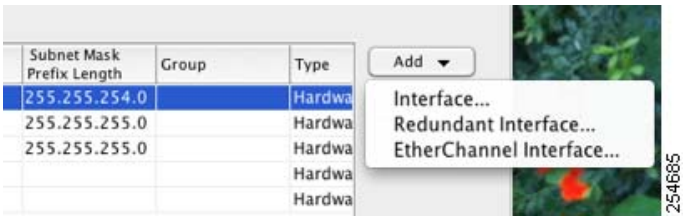

The Add Redundant Interface dialog box appears.

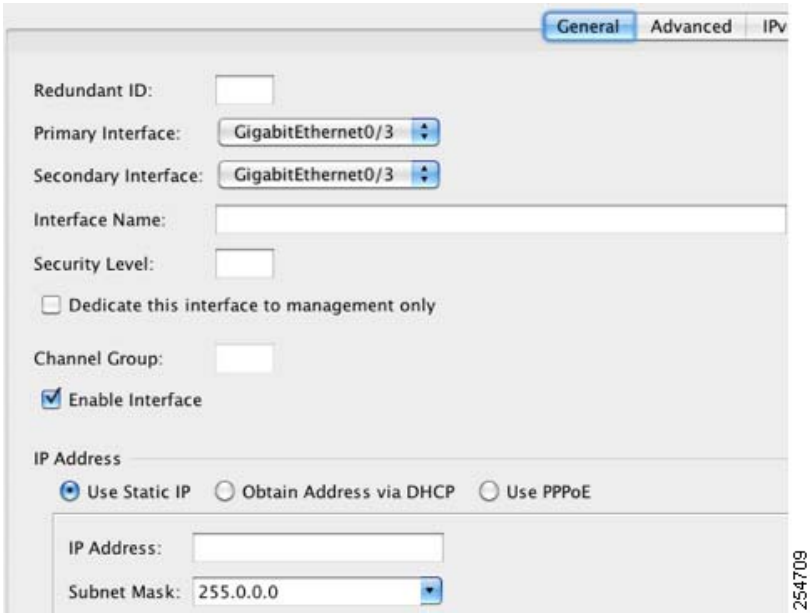

Note In single mode, this procedure only covers a subset of the parameters on the Edit Redundant Interface dialog box; to configure other parameters, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)." Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See the "Configuring Multiple Contexts" section on page 9-15.

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. Redundant interfaces do not support Management *slot*/*port* interfaces as members.

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. The member interfaces now show a lock to the left of the interface ID showing that only basic parameters can be configured for it. The redundant interface is added to the table.

GigabitEthernet0/2	Enabled	No	Redundant8	Hardware	native
GigabitEthernet0/3	Enabled	No		Hardware	native
GigabitEthernet0/3.10	Enabled	No		Logical	vlan100
GigabitEthernet0/3.11	Enabled	No		Logical	vlan11
Management0/0	Enabled	No		Hardware	native
Redundant8	Enabled	Yes		Logical	native

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Configuring Multiple Contexts" section on page 9-15.
- For single context mode, complete the interface configuration. See the Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

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 an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

This section includes the following topics:

- Adding Interfaces to the EtherChannel, page 15-29
- Customizing the EtherChannel, page 15-31

Adding Interfaces to the EtherChannel

This section describes how to create an EtherChannel port-channel interface and assign interfaces to the EtherChannel. By default, port-channel interfaces are enabled.

Guidelines and Limitations

- You can configure up to 48 EtherChannels.
- Each channel group can have eight active interfaces. Note that you can assign up to 16 interfaces to a channel group. While only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- You cannot use interfaces on the 4GE SSM, including the integrated 4GE SSM in slot 1 on the ASA 5550, as part of an EtherChannel.
- To configure a spanned EtherChannel for clustering, see the "Configuring Spanned EtherChannels" section on page 10-37 instead of this procedure.
- See also the "EtherChannel Guidelines" section on page 15-11.

Prerequisites

- All interfaces in the channel group must be the same type, speed, and duplex. Half duplex is not supported.
- You cannot add a physical interface to the channel group 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

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 Choose **Add > EtherChannel Interface**.

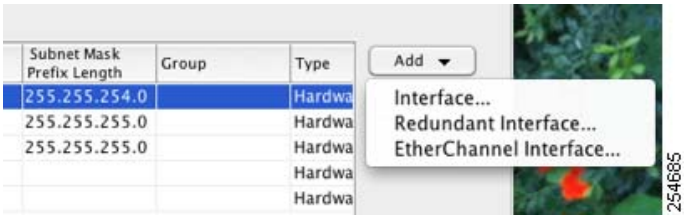

The Add EtherChannel Interface dialog box appears.

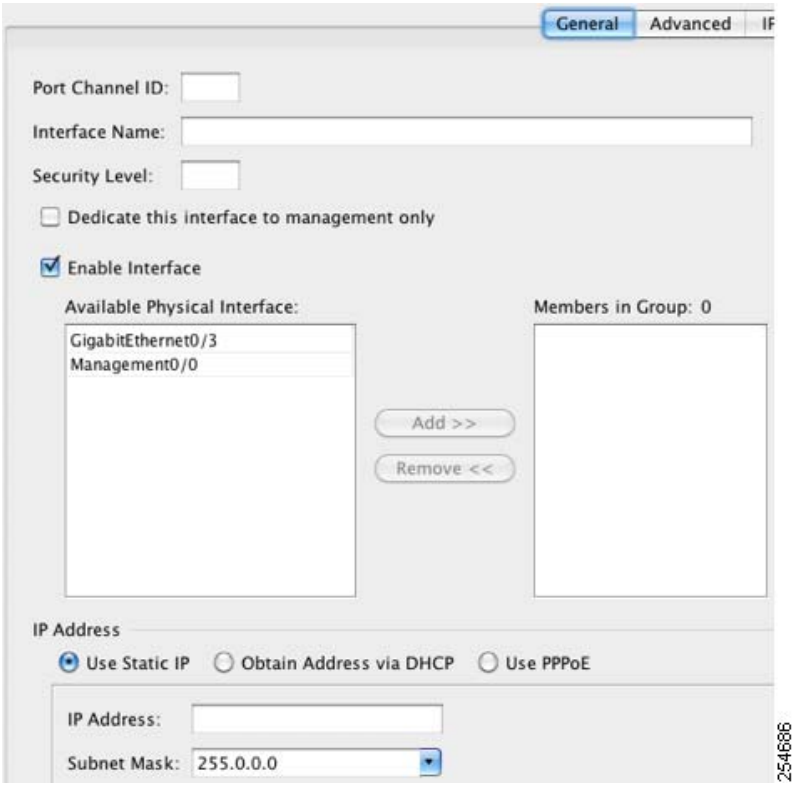

Note In single mode, this procedure only covers a subset of the parameters on the Edit EtherChannel Interface dialog box; to configure other parameters, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)." Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See the "Configuring Multiple Contexts" section on page 9-15.

Step 3 In the Port Channel ID field, enter a number between 1 and 48.

Step 4 In the Available Physical Interface area, click an interface and then click **Add >>** to move it to the Members in Group area.

In transparent mode, if you create a channel group with multiple Management interfaces, then you can use this EtherChannel as the management-only interface.

Step 5 Repeat for each interface you want to add to the channel group.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

Step 6 Click **OK**.

You return to the Interfaces pane. The member interfaces now show a lock to the left of the interface ID showing that only basic parameters can be configured for it. The EtherChannel interface is added to the table.

What to Do Next

Optional Tasks:

- Customize the EtherChannel interface. See the "Customizing the EtherChannel" section on page 15-31.
- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Configuring Multiple Contexts" section on page 9-15.
- For single context mode, complete the interface configuration. See the Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Customizing the EtherChannel

This section describes how to set the maximum number of interfaces in the EtherChannel, the minimum number of operating interfaces for the EtherChannel to be active, the load balancing algorithm, and other optional parameters.

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.

Step 2 Click the port-channel interface you want to customize, and click **Edit**.

The Edit Interface dialog box appears.

Step 3 To override the media type, duplex, speed, and pause frames for flow control for all member interfaces, click **Configure Hardware Properties**. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

a. For slot 1 on the ASA 5550 ASA or the 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** or **Auto**, depending on the interface type, from the Duplex drop-down list. Half is not supported for the EtherChannel.

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 ASA, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. See the "Auto-MDI/MDIX Feature" section on page 15-2.

d. To enable pause (XOFF) frames for flow control on 1-Gigabit and 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 4 To customize the EtherChannel, click the **Advanced** tab.

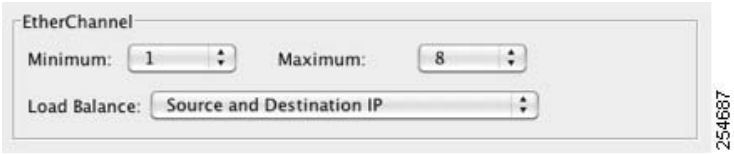

a. In the EtherChannel area, from the Minimum drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 8. The default is 1.

b. From the Maximum drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 8. The default is 8.

c. From the Load Balance drop-down list, select the criteria used to load balance the packets across the group channel interfaces. By default, the ASA balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see the "Load Balancing" section on page 15-7.

Step 5 Click **OK**.

You return to the Interfaces pane.

Step 6 To set the mode and priority for a physical interface in the channel group:

a. Click the physical interface in the Interfaces table, and click **Edit**.

The Edit Interface dialog box appears.

b. Click the **Advanced** tab.

c. In the EtherChannel area, from the Mode drop down list, choose **Active**, **Passive**, or **On**. We recommend using Active mode (the default). For information about active, passive, and on modes, see the "Link Aggregation Control Protocol" section on page 15-6.

d. In the LACP Port Priority field, set the port priority between 1 and 65535. The default is 32768. The higher the number, the lower the priority. The ASA uses this setting to decide which interfaces are active and which are standby if you assign more interfaces than can be used. If the port priority setting is the same for all interfaces, then the priority is determined by the interface ID (slot/port). The lowest interface ID is the highest priority. For example, GigabitEthernet 0/0 is a higher priority than GigabitEthernet 0/1.

If you want to prioritize an interface to be active even though it has a higher interface ID, then set this command to have a lower value. For example, to make GigabitEthernet 1/3 active before GigabitEthernet 0/7, then make the priority value be 12345 on the 1/3 interface vs. the default 32768 on the 0/7 interface.

If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See Step 9 to set the system priority.

Step 7 Click **OK**.

You return to the Interfaces pane.

Step 8 Click **Apply**.

Step 9 To set the LACP system priority, perform the following steps. If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See Step 6d for more information.

a. Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > EtherChannel** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > EtherChannel** pane.

b. In the LACP System Priority field, enter a priority between 1 and 65535.

The default is 32768.

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 15-34.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Configuring Multiple Contexts" section on page 9-15.
- For single context mode, complete the interface configuration. See the Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Configuring VLAN Subinterfaces and 802.1Q Trunking

Subinterfaces let you divide a physical, redundant, or EtherChannel 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 ASAs. 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 ASA 5510 and Higher Interfaces" section on page 15-8.
- 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 Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)," for more information about completing the interface configuration.

- (ASA 5512-X through ASA 5555-X) You cannot configure subinterfaces on the Management 0/0 interface.

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.

Step 2 Choose **Add > Interface**.

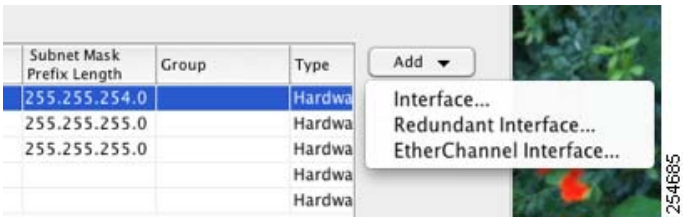

The Add Interface dialog box appears.

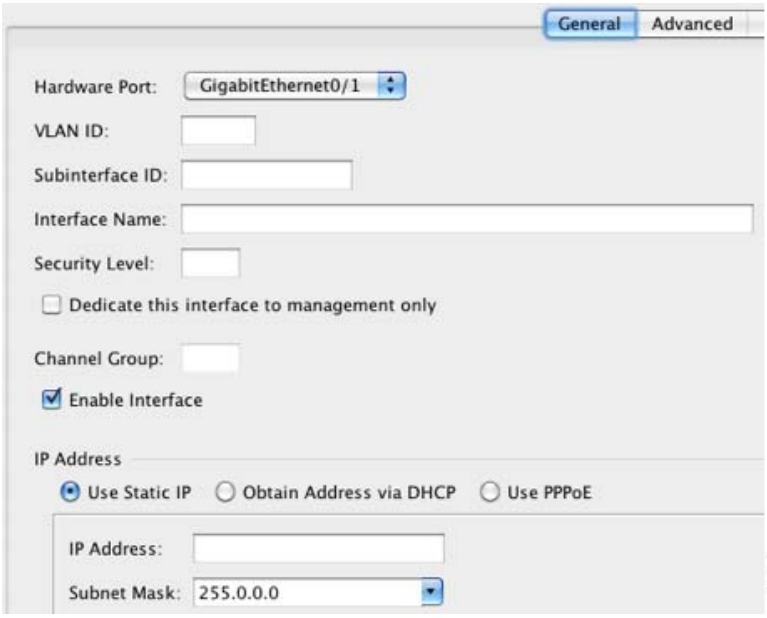

Note In single mode, this procedure only covers a subset of the parameters on the Edit Interface dialog box; to configure other parameters, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)." Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See the "Configuring Multiple Contexts" section on page 9-15.

Step 3 From the Hardware Port drop-down list, choose the physical, redundant, or port-channel 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

(Optional) Enable jumbo frame support according to the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37.

Enabling Jumbo Frame Support (Supported Models)

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.

Supported models include:

- ASA 5512-X
- ASA 5515-X
- ASA 5525-X
- ASA 5545-X
- ASA 5555-X
- ASA 5580
- ASA 5585-X

Prerequisites

- In multiple context mode, set this option in the system execution space.
- Changes in this setting require you to reload the ASA.
- 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 the MAC Address and MTU" section on page 17-13. In multiple context mode, set the MTU within each context.

Detailed Steps

- Multiple mode—To enable jumbo frame support, choose **Configuration > Context Management > Interfaces**, and click the **Enable jumbo frame support** check box.
- Single mode—Setting the MTU larger than 1500 bytes automatically enables jumbo frames. To manually enable or disable this setting, choose **Configuration > Device Setup > Interfaces**, and click the **Enable jumbo frame support** check box.

Monitoring Interfaces

This section includes the following topics:

- ARP Table, page 15-38
- MAC Address Table, page 15-38

- Interface Graphs, page 15-38

ARP Table

The Monitoring > Interfaces > 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.

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 ASA and updates Last Updated date and time.
- Last Updated—*Display only.* Shows the date and time the display was updated.

MAC Address Table

The Monitoring > Interfaces > MAC Address Table pane shows the static and dynamic MAC address entries. See the "MAC Address Table" section on page 15-38 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 the "MAC Address Table" section on page 15-38.
- Refresh—Refreshes the table with current information from the ASA.

Interface Graphs

The Monitoring > Interfaces > Interface Graphs pane lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA 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 ASA was incapable of handing received data to a hardware buffer because the input rate exceeded the ASA capability to handle the data.

  Underruns—The number of times that the transmitter ran faster than the ASA 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 ASA 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 ASA is partly finished sending the packet. The ASA 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.

- 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.

Graph/Table

The Monitoring > Interfaces > Interface Graphs > Graph/Table 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 (see the “Enabling History Metrics” section on page 4-34), 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 (see the “Enabling History Metrics” section on page 4-34). 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 check 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.

Where to Go Next

- For multiple context mode:
  a. Assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Chapter 9, "Configuring Multiple Context Mode."
  b. Complete the interface configuration according to Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."
- For single context mode, complete the interface configuration according to Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Feature History for ASA 5510 and Higher Interfaces

Table 15-3 lists the release history for this feature.

Table 15-3 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 interfaces for the Base license on the ASA 5510	7.2(2)	For the Base license on the ASA 5510, the maximum number of interfaces was increased from 3 plus a management interface to unlimited interfaces.
Increased VLANs	7.2(2)	VLAN limits were increased for the ASA 5510 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).

Table 15-3 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 ASA 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.
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames. 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. This feature is also supported on the ASA 5585-X. We modified the following screen: 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. This feature is also supported on the ASA 5585-X. We modified the following screens: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Support for Pause Frames for Flow Control on 1-Gigabit Ethernet Interfaces	8.2(5)/8.4(2)	You can now enable pause (XOFF) frames for flow control for 1-Gigabit interfaces on all models. We modified the following screens: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
EtherChannel support	8.4(1)	You can configure up to 48 802.3ad EtherChannels of eight active interfaces each. We modified or introduced the following screens: Configuration > Device Setup > Interfaces Configuration > Device Setup > Interfaces > Add/Edit EtherChannel Interface Configuration > Device Setup > Interfaces > Add/Edit Interface Configuration > Device Setup > EtherChannel **Note** EtherChannel is not supported on the ASA 5505.

Starting Interface Configuration (ASA 5505)

This chapter includes tasks for starting your interface configuration for the ASA 5505, including creating VLAN interfaces and assigning them to switch ports.

For ASA 5510 and higher configuration, see the "Feature History for ASA 5505 Interfaces" section on page 16-16.

This chapter includes the following sections:

- Information About ASA 5505 Interfaces, page 16-1
- Licensing Requirements for ASA 5505 Interfaces, page 16-4
- Guidelines and Limitations, page 16-5
- Default Settings, page 16-5
- Starting ASA 5505 Interface Configuration, page 16-6
- Monitoring Interfaces, page 16-12
- Where to Go Next, page 16-15
- Feature History for ASA 5505 Interfaces, page 16-16

Information About ASA 5505 Interfaces

This section describes the ports and interfaces of the ASA 5505 and includes the following topics:

- Understanding ASA 5505 Ports and Interfaces, page 16-2
- Maximum Active VLAN Interfaces for Your License, page 16-2
- VLAN MAC Addresses, page 16-4
- Power over Ethernet, page 16-4
- Monitoring Traffic Using SPAN, page 16-4
- Auto-MDI/MDIX Feature, page 16-4

Understanding ASA 5505 Ports and Interfaces

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

- Physical switch ports—The ASA 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 16-4 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 ASA applies the security policy to the traffic and routes or bridges between the two VLANs.

Maximum Active VLAN Interfaces for Your License

In routed mode, you can configure the following VLANs depending on your license:

- Base license—3 active VLANs. The third VLAN can only be configured to initiate traffic to one other VLAN. See Figure 16-1 for more information.
- Security Plus license—20 active VLANs.

In transparent firewall mode, you can configure the following VLANs depending on your license:

- Base license—2 active VLANs in 1 bridge group.
- Security Plus license—3 active VLANs: 2 active VLANs in 1 bridge group, and 1 active VLAN for the failover link.

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

With the Base license in routed mode, the third VLAN can only be configured to initiate traffic to one other VLAN. See Figure 16-1 for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 16-1 ASA 5505 with Base License

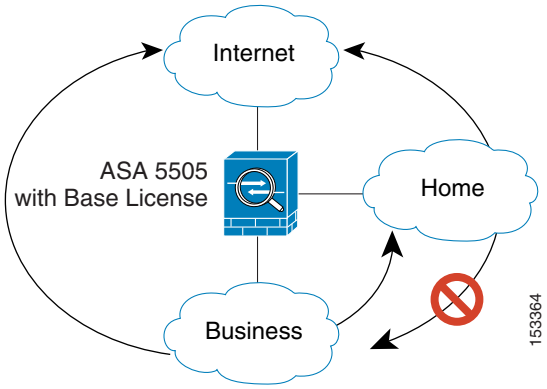

With the Security Plus license, you can configure 20 VLAN interfaces in routed mode, including a VLAN interface for failover and a VLAN interface as a backup link to your ISP. You can configure the backup interface to not pass through traffic unless the route through the primary interface fails. You can configure trunk ports to accommodate multiple VLANs per port.

Note The ASA 5505 supports Active/Standby failover, but not Stateful Failover.

See Figure 16-2 for an example network.

Figure 16-2 ASA 5505 with Security Plus License

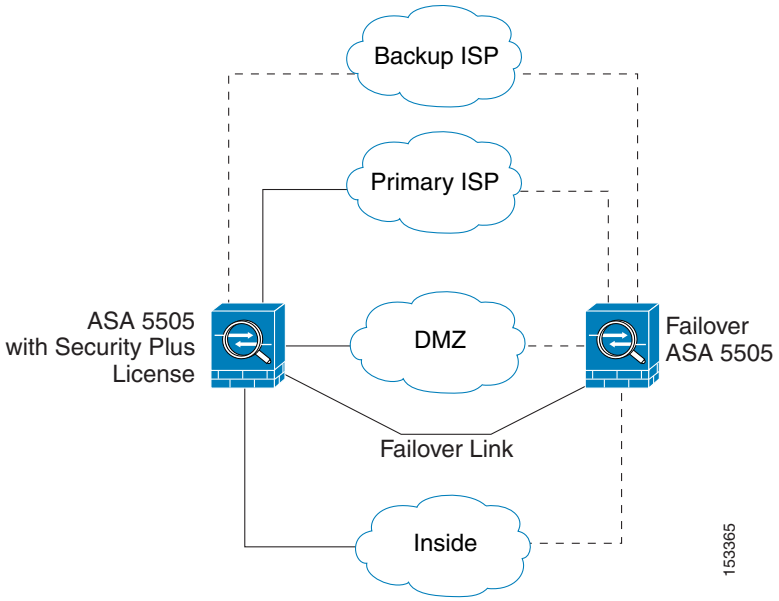

VLAN MAC Addresses

- 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 the MAC Address and MTU" section on page 17-13.
- 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 the MAC Address and MTU" section on page 18-14.

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 ASA 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 16-8 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 command reference for more information.

Auto-MDI/MDIX Feature

All ASA 5505 interfaces include 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. You cannot disable Auto-MDI/MDIX.

Licensing Requirements for ASA 5505 Interfaces

Model	License Requirement
ASA 5505	VLANs: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 VLAN Trunks: Base License: None. Security Plus License: 8. Interfaces of all types[1]: Base License: 52. Security Plus License: 120.

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, and bridge group interfaces.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

The ASA 5505 does not support multiple context mode.

Firewall Mode Guidelines

- In transparent mode, you can configure up to eight bridge groups. Note that you must use at least one bridge group; data interfaces must belong to a bridge group.
- Each bridge group can include up to four VLAN interfaces, up to the license limit.

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 3-17.

Default State of Interfaces

Interfaces have the following default states:

- Switch ports—Disabled.
- VLANs—Enabled. However, for traffic to pass through the VLAN, the switch port must also be enabled.

Default Speed and Duplex

By default, the speed and duplex are set to auto-negotiate.

Starting ASA 5505 Interface Configuration

This section includes the following topics:

- Task Flow for Starting Interface Configuration, page 16-6
- Configuring VLAN Interfaces, page 16-6
- Configuring and Enabling Switch Ports as Access Ports, page 16-8
- Configuring and Enabling Switch Ports as Trunk Ports, page 16-10

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 16-6.

Step 2 Configure and enable switch ports as access ports. See the "Configuring and Enabling Switch Ports as Access Ports" section on page 16-8.

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 16-10.

Step 4 Complete the interface configuration according to Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Configuring VLAN Interfaces

This section describes how to configure VLAN interfaces. For more information about ASA 5505 interfaces, see the "Information About ASA 5505 Interfaces" section on page 16-1.

Guidelines

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

Detailed Steps

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.

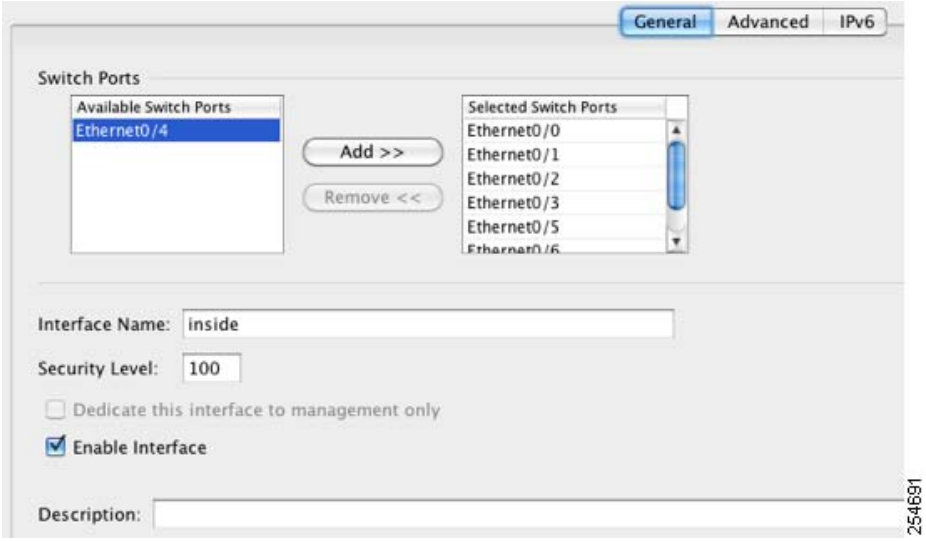

Step 3 In the Available Switch Ports pane, choose a switch port, and click **Add**.

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 Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

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 ASA does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505.

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 the MAC Address and MTU" section on page 17-13.

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 16-8 and the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 16-10.

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 16-10. If you have a factory default configuration, see the "ASA 5505 Default Configuration" section on page 3-20 to check if you want to change the default interface settings according to this procedure.

For more information about ASA 5505 interfaces, see the "Information About ASA 5505 Interfaces" section on page 16-1.

Caution The ASA 5505 does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the ASA 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.

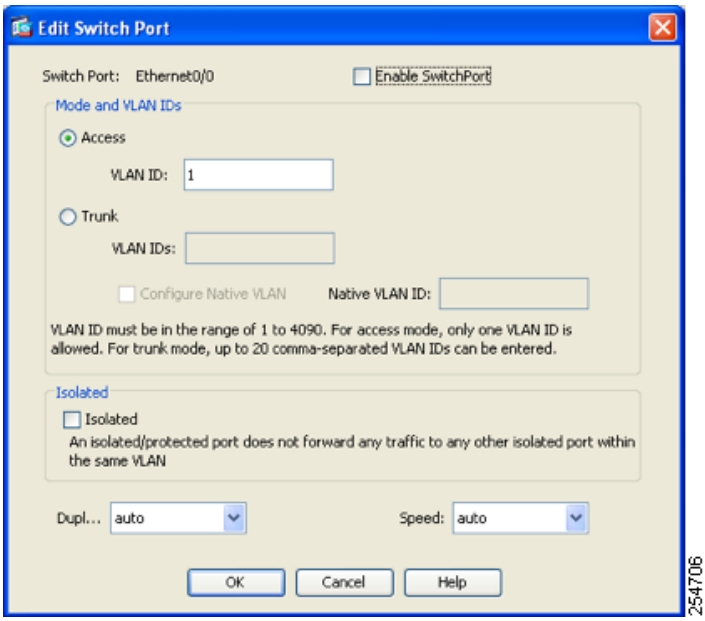

Step 4 To enable the switch port, check the **Enable SwitchPort** check box.

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 16-6 (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 16-6 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 16-6 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 16-10.
- To complete the interface configuration, see Chapter 17, "Completing Interface Configuration (Routed Mode)," or Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Configuring and Enabling Switch Ports as Trunk Ports

This procedure describes 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 16-8.

Guidelines

This switch port cannot pass traffic until you assign at least one VLAN to it, native or non-native.

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.

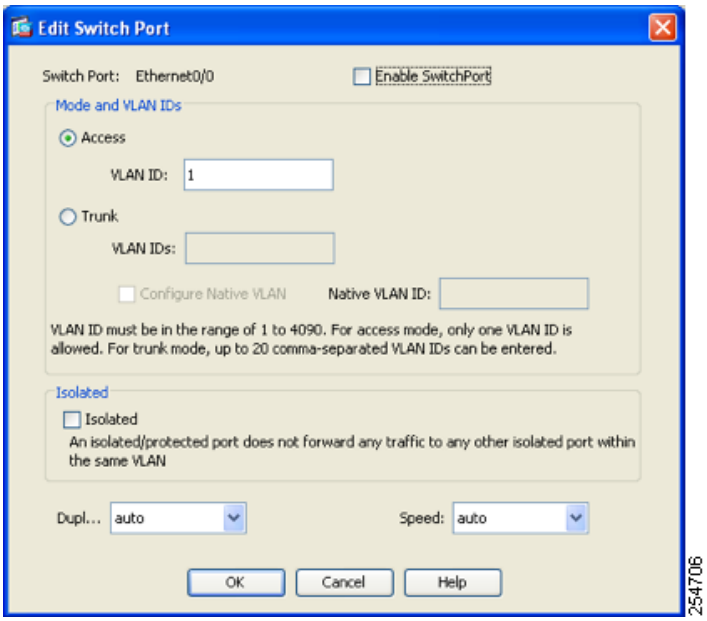

Step 4 To enable the switch port, check the **Enable SwitchPort** check box.

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 16-6 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 16-6 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**.

Monitoring Interfaces

This section includes the following topics:

- ARP Table, page 16-12
- MAC Address Table, page 16-12
- Interface Graphs, page 16-13

ARP Table

The Monitoring > Interfaces > 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.

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 ASA and updates Last Updated date and time.
- Last Updated—*Display only.* Shows the date and time the display was updated.

MAC Address Table

The Monitoring > Interfaces > MAC Address Table pane shows the static and dynamic MAC address entries. See the "MAC Address Table" section on page 16-12 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 the "MAC Address Table" section on page 16-12.
- Refresh—Refreshes the table with current information from the ASA.

Interface Graphs

The Monitoring > Interfaces > Interface Graphs pane lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA 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 ASA was incapable of handing received data to a hardware buffer because the input rate exceeded the ASA capability to handle the data.

    Underruns—The number of times that the transmitter ran faster than the ASA 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 ASA 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 ASA is partly finished sending the packet. The ASA 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.

- 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.

Graph/Table

The Monitoring > Interfaces > Interface Graphs > Graph/Table 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 (see the “Enabling History Metrics” section on page 4-34), 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 (see the “Enabling History Metrics” section on page 4-34). 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 check 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.

Where to Go Next

Complete the interface configuration according to Chapter 17, “Completing Interface Configuration (Routed Mode),” or Chapter 18, “Completing Interface Configuration (Transparent Mode, 8.4 and Later).”

Feature History for ASA 5505 Interfaces

Table 16-1 lists the release history for this feature.

Table 16-1 Feature History for Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.2(2)	The maximum number of VLANs for the Security Plus license on the ASA 5505 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.
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. We modified the following screen: Configuration > Device Setup > Interfaces > Switch Ports > Edit Switch Port.

Completing Interface Configuration (Routed Mode)

This chapter includes tasks to complete the interface configuration for all models in routed firewall mode. This chapter includes the following sections:

- Information About Completing Interface Configuration in Routed Mode, page 17-1
- Licensing Requirements for Completing Interface Configuration in Routed Mode, page 17-2
- Guidelines and Limitations, page 17-5
- Default Settings, page 17-6
- Completing Interface Configuration in Routed Mode, page 17-6
- Turning Off and Turning On Interfaces, page 17-22
- Monitoring Interfaces, page 17-22
- Feature History for Interfaces in Routed Mode, page 17-30

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.

Information About Completing Interface Configuration in Routed Mode

This section includes the following topics:

- Security Levels, page 17-1
- Dual IP Stack (IPv4 and IPv6), page 17-2

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 17-20 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 17-20), 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 ASA.

- 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.

- **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 (IPv4 and IPv6)

The ASA 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.

Licensing Requirements for Completing Interface Configuration in Routed Mode

Model	License Requirement
ASA 5505	VLANs: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 VLAN Trunks: Base License: None. Security Plus License: 8. Interfaces of all types[1]: Base License: 52. Security Plus License: 120.

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, and bridge group interfaces.

Model	License Requirement
ASA 5510	VLANs: Base License: 50 Security Plus License: 100 Interface Speed: Base License—All interfaces Fast Ethernet. Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet. Interfaces of all types[1]: Base License: 52 Security Plus License: 120
ASA 5520	VLANs: Base License: 150. Interfaces of all types[1]: Base License: 640
ASA 5540	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 840
ASA 5550	VLANs: Base License: 400 Interfaces of all types[1]: Base License: 1640

Model	License Requirement
ASA 5580	VLANs: Base License: 1024 Interfaces of all types[1]: Base License: 4176
ASA 5512-X	VLANs: Base License: 50 Interfaces of all types[1]: Base License: 328
ASA 5515-X	VLANs: Base License: 100 Interfaces of all types[1]: Base License: 528
ASA 5525-X	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 928
ASA 5545-X	VLANs: Base License: 300 Interfaces of all types[1]: Base License: 1328
ASA 5555-X	VLANs: Base License: 500 Interfaces of all types[1]: Base License: 2128
ASA 5585-X	VLANs: Base License: 1024 Interface Speed for SSP-10 and SSP-20: Base License—1-Gigabit Ethernet for fiber interfaces 10 GE I/O License (Security Plus)—10-Gigabit Ethernet for fiber interfaces (SSP-40 and SSP-60 support 10-Gigabit Ethernet by default.) Interfaces of all types[1]: Base License: 4176

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces.

Model	License Requirement
ASASM	VLANs: Base License: 1000

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- For the ASA 5510 and higher in multiple context mode, configure the physical interfaces in the system execution space according to Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)." Then, configure the logical interface parameters in the context execution space according to this chapter. For the ASASM in multiple context mode, configure switch ports and VLANs on the switch, and then assign VLANs to the ASASM according to Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

  The ASA 5505 does not support multiple context mode.
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" section on page 9-15.
- PPPoE is not supported in multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode. For transparent mode, see Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See the "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links. In multiple context mode, failover interfaces are configured in the system configuration.

IPv6 Guidelines

Supports IPv6.

VLAN ID Guidelines for the ASASM

You can add any VLAN ID to the configuration, but only VLANs that are assigned to the ASA by the switch can pass traffic. To view all VLANs assigned to the ASA, use the **show vlan** command.

If you add an interface for a VLAN that is not yet assigned to the ASA by the switch, the interface will be in the down state. When you assign the VLAN to the ASA, the interface changes to an up state. See the **show interface** command for more information about interface states.

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 3-17.

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 ASA 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 for the ASASM

- In single mode or in the system execution space, VLAN interfaces are enabled by default.
- 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.

Jumbo Frame Support

By default, the ASASM supports jumbo frames. Just configure the MTU for the desired packet size according to the "Configuring the MAC Address and MTU" section on page 17-13.

Completing Interface Configuration in Routed Mode

This section includes the following topics:

- Task Flow for Completing Interface Configuration, page 17-7
- Configuring General Interface Parameters, page 17-7
- Configuring the MAC Address and MTU, page 17-13
- Configuring IPv6 Addressing, page 17-15
- Allowing Same Security Level Communication, page 17-20

Task Flow for Completing Interface Configuration

Step 1 Set up your interfaces depending on your model:

- ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
- ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
- ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

Step 2 (Multiple context mode) Allocate interfaces to the context according to the "Configuring Multiple Contexts" section on page 9-15.

Step 3 (Multiple context mode) In the Configuration > Device List pane, double-click the context name under the active device IP address.

Step 4 Configure general interface parameters, including the interface name, security level, and IPv4 address. See the "Configuring General Interface Parameters" section on page 17-7.

Step 5 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 17-13.

Step 6 (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 17-15.

Step 7 (Optional) Allow same security level communication, either by allowing communication between two interfaces or by allowing traffic to enter and exit the same interface. See the "Allowing Same Security Level Communication" section on page 17-20.

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
- EtherChannel interfaces

For the ASA 5505 and ASASM, you must configure interface parameters for the following interface types:

- VLAN interfaces

Guidelines and Limitations

- For the ASA 5550, 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.
- 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 "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links.

Restrictions

- PPPoE is not supported in multiple context mode.
- PPPoE and DHCP are not supported on the ASASM.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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.

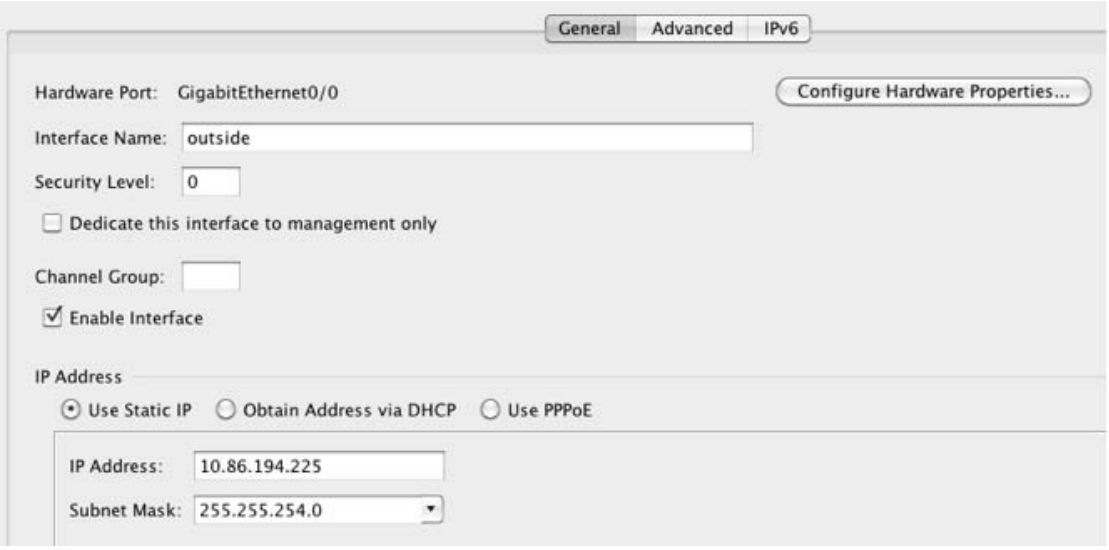

Step 3 In the Interface Name field, enter a name up to 48 characters in length.

Step 4 In the Security level field, enter a level between 0 (lowest) and 100 (highest).

See the "Security Levels" section on page 17-1 for more information.

Step 5 (Optional; not supported for redundant interfaces) 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 "Prerequisites" section on page 17-8 for more information.

(ASA 5512-X through ASA 5555-X) You cannot disable this option on the Management 0/0 interface.

Note The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

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.

- 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.

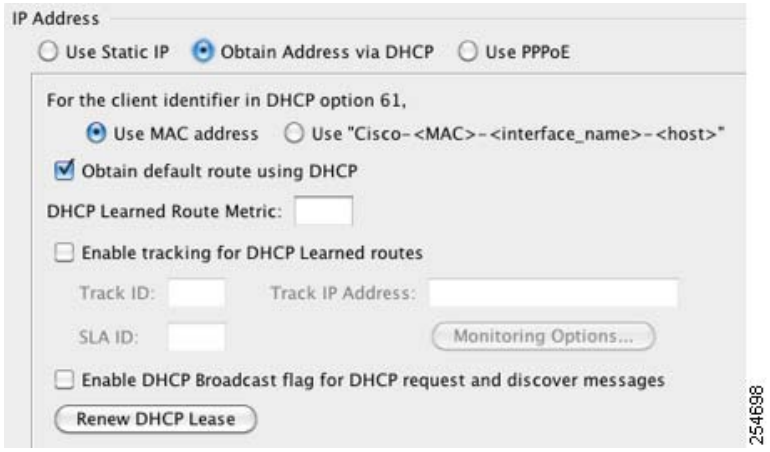

a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

   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. To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.

c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.

d. (Optional) To assign 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.

e. (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.

f. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

g. (Optional) To renew the lease, click **Renew DHCP Lease**.

- (Single mode only) To obtain an IP address using PPPoE, check **Use PPPoE**.

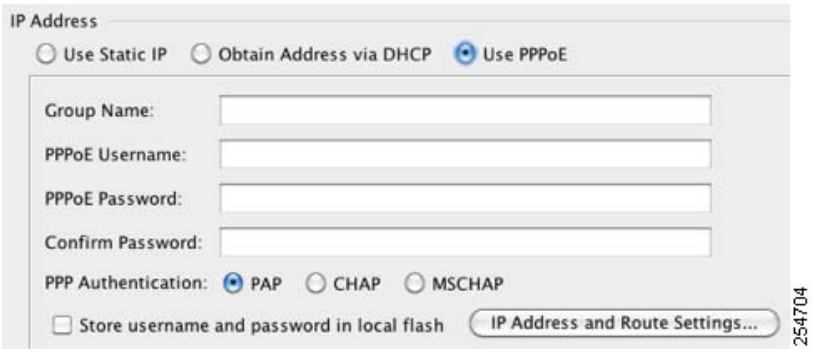

a. In the Group Name field, specify a group name.

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 ASA stores the username and password in a special location of NVRAM. If an Auto Update Server sends a **clear configure** command to the ASA, and the connection is then interrupted, the ASA 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 17-11 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 15-23.

Step 9 Click **OK**.

What to Do Next

- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 17-13.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 17-15.

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.

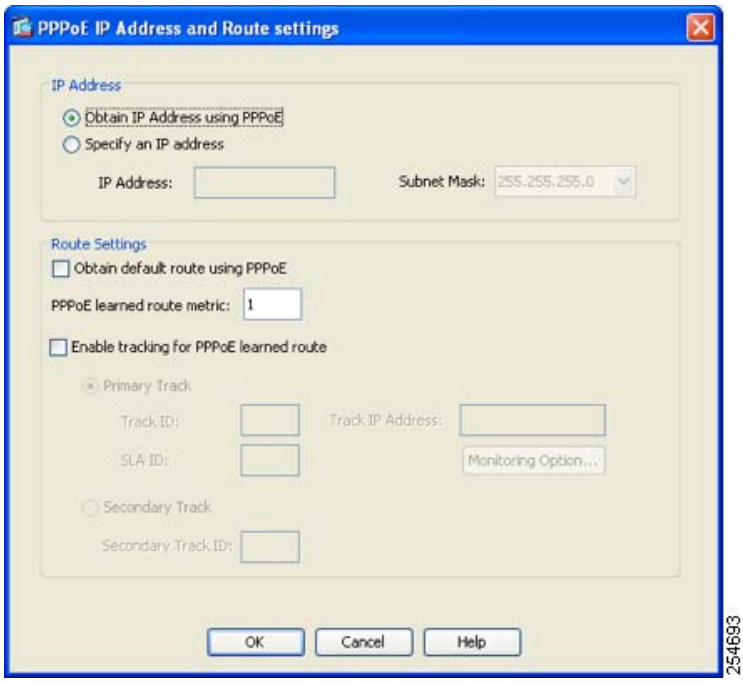

Fields

- 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 ASA to use PPP to get an IP address.
  - Specify an IP Address—Specify an IP address and mask for the ASA 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 the MAC Address and MTU

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.

For the ASASM, all VLANs use the same MAC address provided by the backplane.

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.

For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. In multiple context mode, you can automatically assign unique MAC addresses to interfaces, including an EtherChannel port interface. We recommend manually, or in multiple context mode, automatically configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

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 ASA 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 ASA Classifies Packets" section on page 9-3 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 to Context Interfaces" section on page 9-24 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.

Information About the MTU

The MTU is the maximum datagram size that is sent on a connection. Data that is larger than the MTU value is fragmented before being sent.

The ASA supports IP path MTU discovery (as defined in RFC 1191), which allows a host to dynamically discover and cope with the differences in the maximum allowable MTU size of the various links along the path. Sometimes, the ASA cannot forward a datagram because the packet is larger than the MTU that you set for the interface, but the "don't fragment" (DF) bit is set. The network software sends a message to the sending host, alerting it to the problem. The host has to fragment packets for the destination so that they fit the smallest packet size of all the links along the path.

The default MTU is 1500 bytes in a block for Ethernet interfaces. This value is sufficient for most applications, but you can pick a lower number if network conditions require it.

Jumbo frames are supported by default on the ASASM. To enable jumbo frames, see the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37. 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. To use jumbo frames, set the value higher, for example, to 9000 bytes.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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 Click the **Advanced** tab.

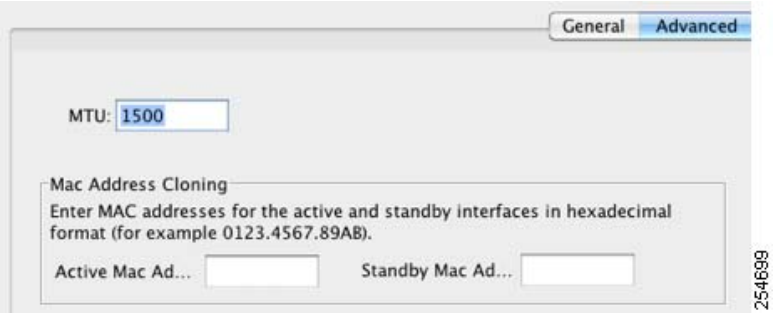

Step 4 To set the MTU or to enable jumbo frame support (supported models only), enter the value in the MTU field, between 300 and 65,535 bytes.

The default is 1500 bytes.

Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.

- For models that support jumbo frames 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 models that support jumbo frames 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 (Supported Models)" section on page 15-37.

Note Enabling or disabling jumbo frame support requires you to reload the ASA.

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. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

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 17-15.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the "IPv6 Addresses" section on page A-5.

This section includes the following topics:

- Information About IPv6, page 17-16

- Configuring a Global IPv6 Address, page 17-16
- Configuring IPv6 Neighbor Discovery, page 17-18
- (Optional) Configuring the Link-Local Addresses Automatically, page 17-19
- (Optional) Configuring the Link-Local Addresses Manually, page 17-19

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 17-16
- Modified EUI-64 Interface IDs, page 17-16

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network.
- 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.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

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 ASA can enforce this requirement for hosts attached to the local link.

When this feature 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:

```
%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.

Configuring a Global IPv6 Address

To configure a global IPv6 address, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately.

Restrictions

The ASA does not support IPv6 anycast addresses.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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.

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.

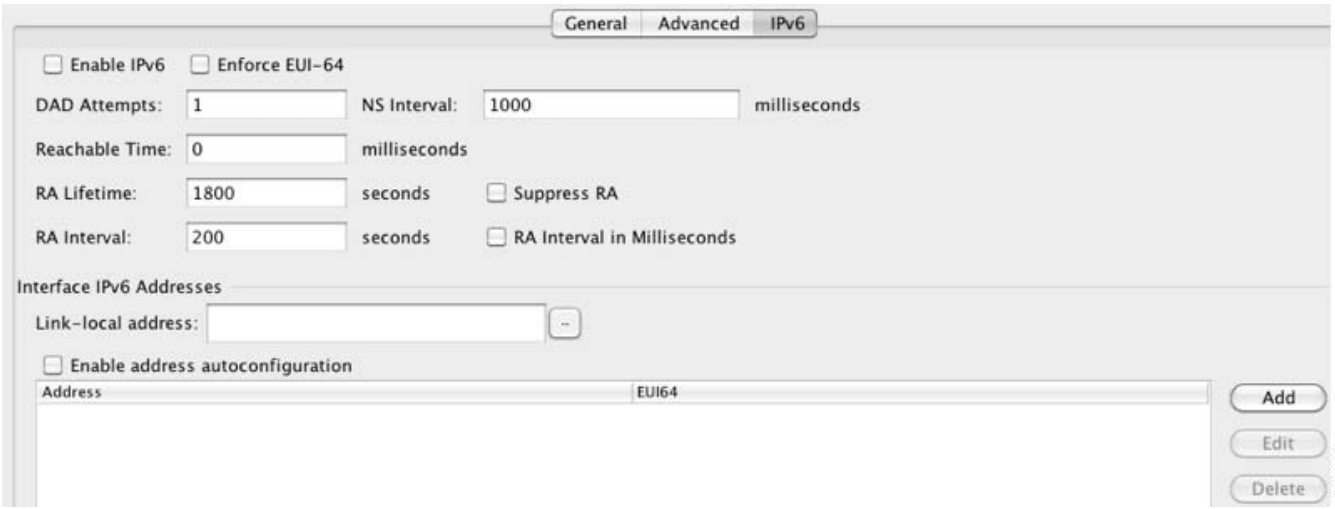

Step 4 Check the **Enable IPv6** check box.

Step 5 (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.

See the "Modified EUI-64 Interface IDs" section on page 17-16 for more information.

Step 6 (Optional) In the top area, customize the IPv6 configuration by referring to Chapter 37, "Configuring IPv6 Neighbor Discovery."

Step 7 Configure the global IPv6 address using one of the following methods.

- 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.

  Note Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the ASA does send Router Advertisement messages in this case. See the **Suppress RA chck box** to suppress messages.

- 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 either full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. If you only enter the prefix, then be sure to check the **EUI 64** checkbox to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked). See the "IPv6 Addresses" section on page A-5 for more information about IPv6 addressing.

  Note For information about the ASA Cluster IP Pool, see the "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-32.

  c. Click **OK**.

Step 8 (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, refer to the "Configuring the IPv6 Prefix in Router Advertisements" section on page 37-11.

Step 9 Click **OK**.

You return to the Configuration > Device Setup > Interfaces pane.

Configuring IPv6 Neighbor Discovery

See Chapter 37, "Configuring IPv6 Neighbor Discovery," to configure IPv6 neighbor discovery.

(Optional) Configuring the Link-Local Addresses Automatically

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required for the interface ID.)

To manually assign the link-local address (not recommended), see the "(Optional) Configuring the Link-Local Addresses Manually" section on page 17-19.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 17-16.

To automatically configure the link-local addresses for 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 In the IPv6 configuration area, check the **Enable IPv6** check box.

This option enables IPv6 and automatically generates the link-local address using the Modified EUI-64 interface ID based on the interface MAC address.

Step 5 Click **OK**.

(Optional) Configuring the Link-Local Addresses Manually

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

To automatically assign the link-local address (recommended), see the "(Optional) Configuring the Link-Local Addresses Automatically" section on page 17-19.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 17-16.

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 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 A-5 for more information about IPv6 addressing.

Step 5 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.

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 ASA is the hub, and remote VPN networks are spokes, for one spoke to communicate with another spoke, traffic must go into the ASA 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 ASA.

For the ASASM, before you can enable this feature, you must first correctly configure the MSFC so that packets are sent to the ASA MAC address instead of being sent directly through the switch to the destination host. Figure 17-1 shows a network where hosts on the same interface need to communicate.

Figure 17-1 Communication Between Hosts on the Same Interface

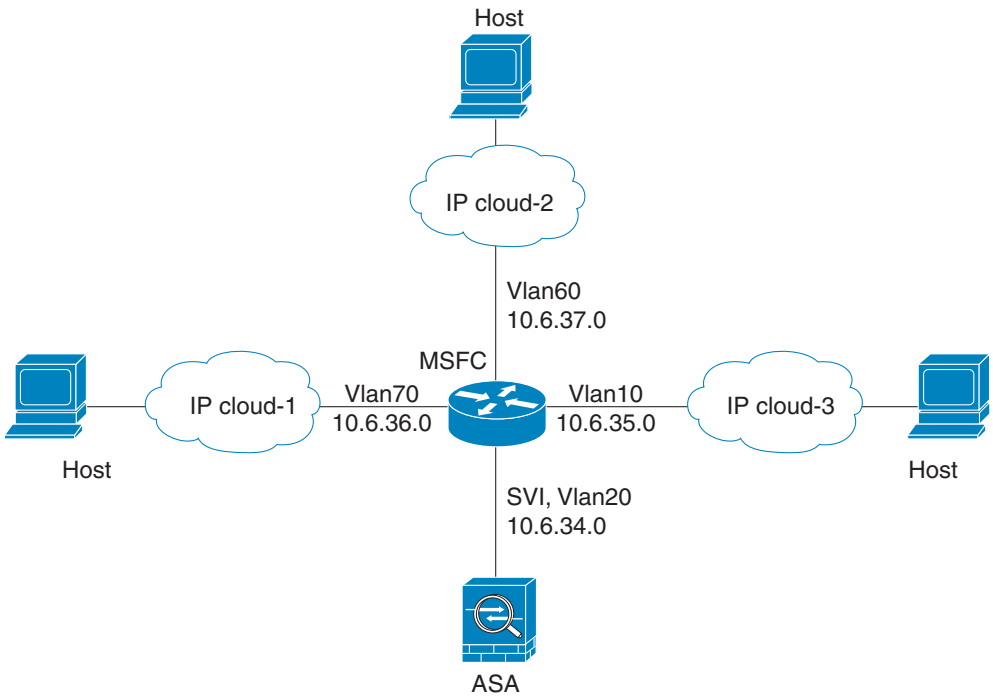

The following sample configuration shows the Cisco IOS **route-map** commands used to enable policy routing in the network shown in Figure 17-1:

```
route-map intra-inter3 permit 0
    match ip address 103
    set interface Vlan20
    set ip next-hop 10.6.34.7
!
route-map intra-inter2 permit 20
    match ip address 102
    set interface Vlan20
    set ip next-hop 10.6.34.7
!
route-map intra-inter1 permit 10
    match ip address 101
    set interface Vlan20
    set ip next-hop 10.6.34.7
```

Detailed Steps

- 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**.

Turning Off and Turning On Interfaces

This section describes how to turn off and on an interface.

All interfaces are enabled by default. In multiple context mode, if you disable or reenable the interface within a context, only that context interface is affected. But if you disable or reenable the interface in the system execution space, then you affect that interface for all contexts.

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 VLAN interface that you want to configure, and click **Edit**.

The Edit Interface dialog box appears.

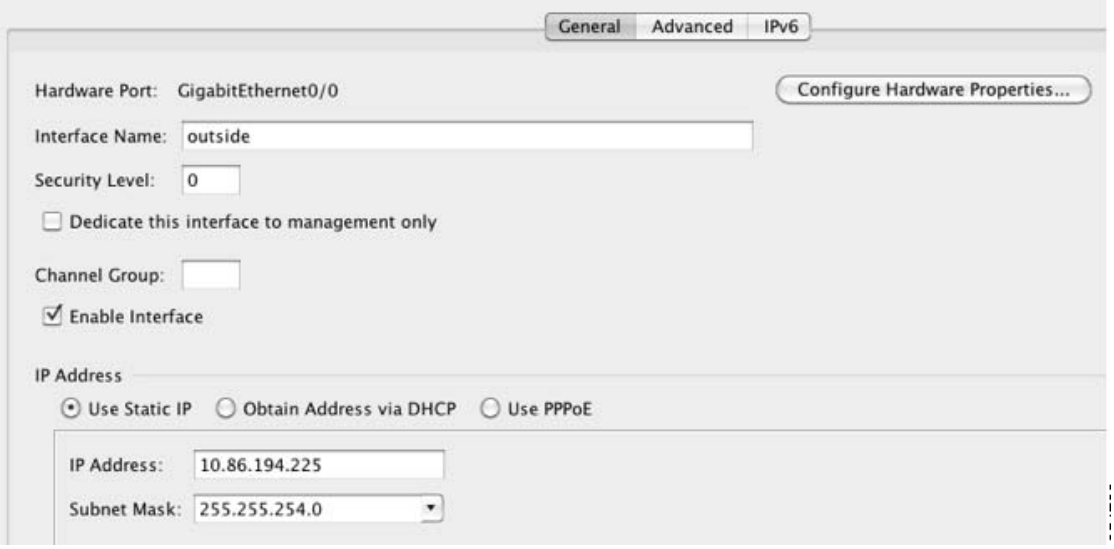

Step 3 To enable or disable the interface, check or uncheck the **Enable Interface** check box.

Monitoring Interfaces

This section includes the following topics:

- ARP Table, page 17-23
- DHCP, page 17-23
- MAC Address Table, page 17-26
- Dynamic ACLs, page 17-26

- Interface Graphs, page 17-26
- PPPoE Client, page 17-29
- Interface Connection, page 17-29

ARP Table

The Monitoring > Interfaces > 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.

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 ASA and updates Last Updated date and time.
- Last Updated—*Display only.* Shows the date and time the display was updated.

DHCP

The ASA lets you monitor DHCP status, including the addresses assigned to clients, the lease information for the ASA interface, and DHCP statistics.

DHCP Server Table

The Monitoring > Interfaces > DHCP > 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 ASA.
- Last Updated—Shows when the data in the table was last updated.

DHCP Client Lease Information

If you obtain the ASA interface IP address from a DHCP server, the Monitoring > Interfaces > DHCP > DHCP Server Table > DHCP Client Lease Information pane shows information about the DHCP lease.

Fields

- Select an interface—Lists the ASA 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 ASA begins the process of acquiring a lease. This state is also shown when a lease ends or when a lease negotiation fails.

    Selecting—The ASA is waiting to receive DHCPOFFER messages from one or more DHCP servers, so it can choose one.

    Requesting—The ASA is waiting to hear back from the server to which it sent its request.

    Purging—The ASA is removing the lease because of an error.

    Bound—The ASA has a valid lease and is operating normally.

    Renewing—The ASA is trying to renew the lease. It regularly sends DHCPREQUEST messages to the current DHCP server, and waits for a reply.

    Rebinding—The ASA 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 ASA started the process to remove the lease.

    Releasing—The ASA 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 ASA attempts to rebind to a DHCP server. Rebinding occurs if the ASA cannot communicate with the original DHCP server, and 87.5 percent of the lease time has expired. The ASA 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 ASA is attempting to establish a lease, this field shows the number of times the ASA tried sending a DHCP message. For example, if the ASA is in the Selecting state, this value shows the number of times the ASA sent discover messages. If the ASA is in the Requesting state, this value shows the number of times the ASA 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.

DHCP Statistics

The Monitoring > Interfaces > DHCP > 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 ASA.
- Total Messages Sent—Shows the total number of messages sent by the ASA.
- 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.

MAC Address Table

The Monitoring > Interfaces > MAC Address Table pane shows the static and dynamic MAC address entries. See the "MAC Address Table" section on page 17-26 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 the "MAC Address Table" section on page 17-26.
- Refresh—Refreshes the table with current information from the ASA.

Dynamic ACLs

The Monitoring > Interfaces > 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 ASA. 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.

Interface Graphs

The Monitoring > Interfaces > Interface Graphs pane lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA 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 ASA was incapable of handing received data to a hardware buffer because the input rate exceeded the ASA capability to handle the data.

  Underruns—The number of times that the transmitter ran faster than the ASA 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 ASA 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 ASA is partly finished sending the packet. The ASA 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.

- 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.

Graph/Table

The Monitoring > Interfaces > Interface Graphs > Graph/Table 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 (see the "Enabling History Metrics" section on page 4-34), 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 (see the "Enabling History Metrics" section on page 4-34). 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 check 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.

PPPoE Client

The Monitoring > Interfaces > PPPoE Client > 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 ASA for display.

Interface Connection

The Monitoring > Interfaces > *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 17-29
- Monitoring Statistics for, page 17-29

Track Status for

The Monitoring > Interfaces > interface connection > Track Status for pane displays information about 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.

Monitoring Statistics for

The Monitoring > Interfaces > interface connection > Monitoring Statistics 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.

Feature History for Interfaces in Routed Mode

Table 17-1 lists the release history for this feature.

Table 17-1 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 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 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 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.
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. We modified the following screen: Configuration > Device Setup > Interfaces > Switch Ports > Edit Switch Port.
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames. 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. We modified the following screen: Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced.

Table 17-1 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
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. We modified the following screens: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface

Completing Interface Configuration (Transparent Mode, 8.4 and Later)

This chapter includes tasks to complete the interface configuration for all models in transparent firewall mode.

For Version 8.3 and earlier, see Chapter 19, "Completing Interface Configuration (Transparent Mode, 8.3 and Earlier)."

This chapter includes the following sections:

- Information About Completing Interface Configuration in Transparent Mode (8.4 and Later), page 18-1
- Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 18-3
- Guidelines and Limitations, page 18-5
- Default Settings, page 18-6
- Completing Interface Configuration in Transparent Mode (8.4 and Later), page 18-7
- Turning Off and Turning On Interfaces, page 18-21
- Monitoring Interfaces, page 18-22
- Feature History for Interfaces in Transparent Mode, page 18-30

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.

Information About Completing Interface Configuration in Transparent Mode (8.4 and Later)

This section includes the following topics:

- Bridge Groups in Transparent Mode, page 18-2
- Security Levels, page 18-2

Bridge Groups in Transparent Mode

If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the ASA, and traffic must exit the ASA before it is routed by an external router back to another bridge group in the ASA. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration. For complete security policy separation, use security contexts with one bridge group in each context. At least one bridge group is required per context or in single mode.

Each bridge group requires a management IP address. For another method of management, see the "Management Interface" section.

Note The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.

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 18-21 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 18-21), 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 ASA.

- 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.

- **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.

Licensing Requirements for Completing Interface Configuration in Transparent Mode

Model	License Requirement
ASA 5505	VLANs: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 VLAN Trunks: Base License: None. Security Plus License: 8. Interfaces of all types[1]: Base License: 52. Security Plus License: 120.

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, and bridge group interfaces.

Model	License Requirement
ASA 5510	VLANs: Base License: 50 Security Plus License: 100 Interface Speed: Base License—All interfaces Fast Ethernet. Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet. Interfaces of all types[1]: Base License: 52 Security Plus License: 120
ASA 5520	VLANs: Base License: 150. Interfaces of all types[1]: Base License: 640
ASA 5540	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 840

Model	License Requirement
ASA 5550	VLANs: Base License: 400 Interfaces of all types[1]: Base License: 1640
ASA 5580	VLANs: Base License: 1024 Interfaces of all types[1]: Base License: 4176
ASA 5512-X	VLANs: Base License: 50 Interfaces of all types[1]: Base License: 328
ASA 5515-X	VLANs: Base License: 100 Interfaces of all types[1]: Base License: 528
ASA 5525-X	VLANs: Base License: 200 Interfaces of all types[1]: Base License: 928
ASA 5545-X	VLANs: Base License: 300 Interfaces of all types[1]: Base License: 1328
ASA 5555-X	VLANs: Base License: 500 Interfaces of all types[1]: Base License: 2128
ASA 5585-X	VLANs: Base License: 1024 Interface Speed for SSP-10 and SSP-20: Base License—1-Gigabit Ethernet for fiber interfaces 10 GE I/O License (Security Plus)—10-Gigabit Ethernet for fiber interfaces (SSP-40 and SSP-60 support 10-Gigabit Ethernet by default.) Interfaces of all types[1]: Base License: 4176

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces.

Model	License Requirement
ASASM	VLANs: Base License: 1000

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- For the ASA 5510 and higher in multiple context mode, configure the physical interfaces in the system execution space according to Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)." Then, configure the logical interface parameters in the context execution space according to this chapter. For the ASASM in multiple context mode, configure switch ports and VLANs on the switch, and then assign VLANs to the ASASM according to Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

  The ASA 5505 does not support multiple context mode.

- You can only configure context interfaces that you already assigned to the context in the system configuration.

Firewall Mode Guidelines

- You can configure up to 8 bridge groups in single mode or per context in multiple mode. Note that you must use at least 1 bridge group; data interfaces must belong to a bridge group.

  Note Although you can configure multiple bridge groups on the ASA 5505, the restriction of 2 data interfaces in transparent mode on the ASA 5505 means you can only effectively use 1 bridge group.

- Each bridge group can include up to 4 interfaces.
- For IPv4, a management IP address is required for each bridge group for both management traffic and for traffic to pass through the ASA.

  Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire bridge group. The ASA uses this IP address as the source address for packets originating on the ASA, such as system messages or AAA communications. In addition to the bridge group management address, you can optionally configure a management interface for some models; see the "Management Interface" section on page 15-2 for more information.

  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). The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported. See the "Configuring Bridge Groups" section on page 18-8 for more information about management IP subnets.

- 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 ASA, you need to configure a global IPv6 address for each bridge group.
- 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.

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See the "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links. In multiple context mode, failover interfaces are configured in the system configuration.

IPv6 Guidelines

- Supports IPv6.
- No support for IPv6 anycast addresses in transparent mode.

VLAN ID Guidelines for the ASASM

You can add any VLAN ID to the configuration, but only VLANs that are assigned to the ASA by the switch can pass traffic. To view all VLANs assigned to the ASA, use the **show vlan** command.

If you add an interface for a VLAN that is not yet assigned to the ASA by the switch, the interface will be in the down state. When you assign the VLAN to the ASA, the interface changes to an up state. See the **show interface** command for more information about interface states.

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 3-17.

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 ASA 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 for the ASASM

- In single mode or in the system execution space, VLAN interfaces are enabled by default.
- 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.

Jumbo Frame Support

By default, the ASASM supports jumbo frames. Just configure the MTU for the desired packet size according to the "Configuring the MAC Address and MTU" section on page 18-14.

Completing Interface Configuration in Transparent Mode (8.4 and Later)

This section includes the following topics:

- Task Flow for Completing Interface Configuration, page 18-7
- Configuring Bridge Groups, page 18-8
- Configuring General Interface Parameters, page 18-9
- Configuring a Management Interface (ASA 5510 and Higher), page 18-11
- Configuring the MAC Address and MTU, page 18-14
- Configuring IPv6 Addressing, page 18-16
- Allowing Same Security Level Communication, page 18-21

Task Flow for Completing Interface Configuration

Step 1 Set up your interfaces depending on your model:

- ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
- ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
- ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

Step 2 (Multiple context mode) Allocate interfaces to the context according to the "Configuring Multiple Contexts" section on page 9-15.

Step 3 (Multiple context mode) In the Configuration > Device List pane, double-click the context name under the active device IP address.

Step 4 Configure one or more bridge groups, including the IPv4 address. See the "Configuring Bridge Groups" section on page 18-8.

Step 5 Configure general interface parameters, including the interface name and security level. See the "Configuring General Interface Parameters" section on page 18-9.

Step 6 (Optional; not supported for the ASA 5505) Configure a management interface. See the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 18-11.

Step 7 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 18-14.

Step 8 (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 18-16.

Step 9 (Optional) Allow same security level communication, either by allowing communication between two interfaces or by allowing traffic to enter and exit the same interface. See the "Allowing Same Security Level Communication" section on page 18-21.

Configuring Bridge Groups

Each bridge group requires a management IP address. The ASA uses this IP address as the source address for packets originating from the bridge group. The management IP address must be on the same subnet as the connected network. 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.

Guidelines and Limitations

You can configure up to 8 bridge groups in single mode or per context in multiple mode. Note that you must use at least one bridge group; data interfaces must belong to a bridge group.

Note For a separate management interface (for supported models), a non-configurable bridge group (ID 101) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Detailed Steps

Step 1 Choose the **Configuration > Interfaces** pane, and choose **Add > Bridge Group Interface**.

The Add Bridge Group dialog box appears.

Step 2 In the Bridge Group ID field, enter the bridge group ID between 1 and 100.

Step 3 In the IP Address field, enter the management IPv4 address.

The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.

Step 4 In the Subnet Mask field, enter the subnet mask or choose one from the menu.

Do not assign a host address (/32 or 255.255.255.255) to the transparent firewall. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The ASA drops all ARP packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the ASA drops the ARP request from the downstream router to the upstream router.

Step 5 (Optional) In the Description field, enter a description for this bridge group.

Step 6 Click **OK**.

Step 7 A Bridge Group Virtual Interface (BVI) is added to the interface table, along with the physical and subinterfaces.

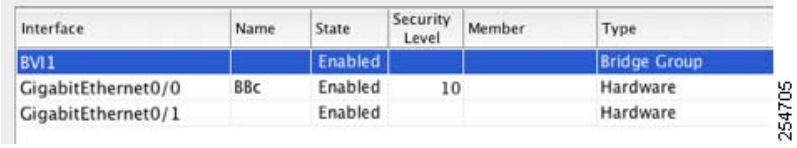

What to Do Next

Configure general interface parameters. See the "Configuring General Interface Parameters" section on page 18-9.

Configuring General Interface Parameters

This procedure describes how to set the name, security level, and bridge group for each transparent interface.

To configure a separate management interface, see the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 18-11.

For the ASA 5510 and higher, you must configure interface parameters for the following interface types:

- Physical interfaces
- VLAN subinterfaces
- Redundant interfaces
- EtherChannel interfaces

For the ASA 5505 and the ASASM, you must configure interface parameters for the following interface types:

- VLAN interfaces

Guidelines and Limitations

- You can configure up to four interfaces per bridge group.
- For the ASA 5550, 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 18-2.
- 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 "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."

- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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.

BVIs appear in the table alongside physical interfaces, subinterfaces, redundant interfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a non-BVI interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

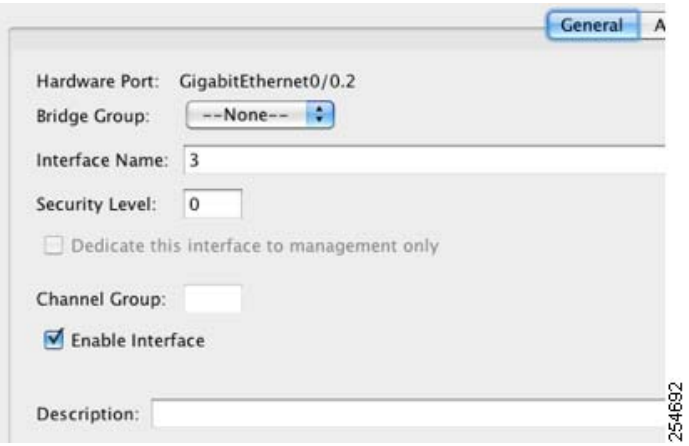

Do not use this procedure for Management interfaces; see the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 18-11 to configure the Management interface.

Step 3 In the Bridge Group drop-down menu, choose the bridge group to which you want to assign this interface.

Step 4 In the Interface Name field, enter a name up to 48 characters in length.

Step 5 In the Security level field, enter a level between 0 (lowest) and 100 (highest).

See the "Security Levels" section on page 18-2 for more information.

Note Do not click the **Dedicate this interface to management only** check box; see the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 18-11 for this option.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Note The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

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. 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, single mode) For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

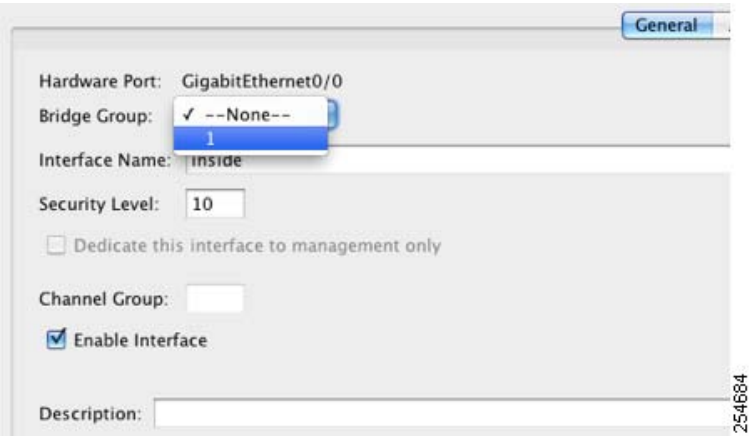

Step 8 Click **OK**.

What to Do Next

- (Optional) Configure a management interface. See the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 18-11.
- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 18-14.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 18-16.

Configuring a Management Interface (ASA 5510 and Higher)

You can configure one management interface separate from the bridge group interfaces in single mode or per context. For more information, see the "Management Interface" section on page 15-2.

Restrictions

- See the "Management Interface" section on page 15-2.
- Do not assign this interface to a bridge group; a non-configurable bridge group (ID 101) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

- If your model does not include a Management interface, you must manage the transparent firewall from a data interface; skip this procedure. (For example, on the ASA 5505.)
- In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. To provide management per context, you can create subinterfaces of the Management interface and allocate a Management subinterface to each context. Note that the ASA 5512-X through ASA 5555-X do not allow subinterfaces on the Management interface, so for per-context management, you must connect to a data interface.

Prerequisites

- Complete the procedures in Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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.

BVIs appear in the table alongside physical interfaces, subinterfaces, redundant interfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a Management interface, subinterface, or EtherChannel port-channel interface comprised of Management interfaces, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

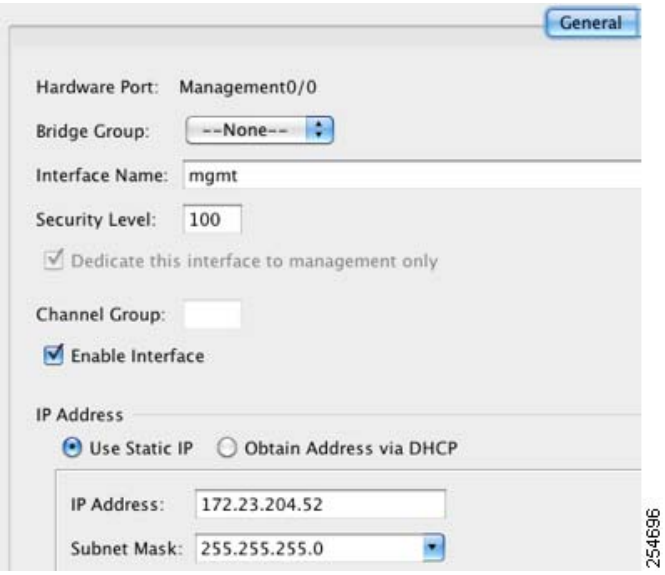

Step 3 In the Bridge Group drop-down menu, leave the default **--None--**. You cannot assign a management interface to a bridge group.

Step 4 In the Interface Name field, enter a name up to 48 characters in length.

Step 5 In the Security level field, enter a level between 0 (lowest) and 100 (highest).

See the "Security Levels" section on page 18-2 for more information.

Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

Note For use with failover, you must set the IP address and standby address manually; DHCP is not supported. Set the standby IP addresses on the Configuration > Device Management > High Availability > Failover > Interfaces tab.

- 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.

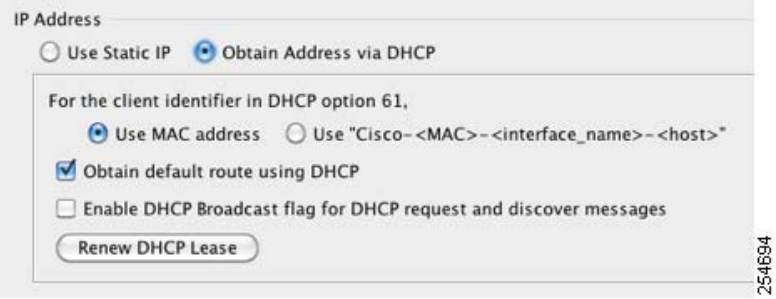

a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

   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. To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.

c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.

d. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

   The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

e. (Optional) To renew the lease, click **Renew DHCP Lease**.

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.

Note (ASA 5510 and higher, single mode) For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Step 9 Click **OK**.

What to Do Next

- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 18-14.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 18-16.

Configuring the MAC Address and MTU

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.

For the ASASM, all VLANs use the same MAC address provided by the backplane.

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.

For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. In multiple context mode, you can automatically assign unique MAC addresses to interfaces, including an EtherChannel port interface. We recommend manually, or in multiple context mode, automatically configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

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 ASA 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 ASA Classifies Packets" section on page 9-3 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 to Context Interfaces" section on page 9-24 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.

Information About the MTU

The MTU is the maximum datagram size that is sent on a connection. Data that is larger than the MTU value is fragmented before being sent.

The ASA supports IP path MTU discovery (as defined in RFC 1191), which allows a host to dynamically discover and cope with the differences in the maximum allowable MTU size of the various links along the path. Sometimes, the ASA cannot forward a datagram because the packet is larger than the MTU that you set for the interface, but the "don't fragment" (DF) bit is set. The network software sends a message to the sending host, alerting it to the problem. The host has to fragment packets for the destination so that they fit the smallest packet size of all the links along the path.

The default MTU is 1500 bytes in a block for Ethernet interfaces. This value is sufficient for most applications, but you can pick a lower number if network conditions require it.

Jumbo frames are supported by default on the ASASM. To enable jumbo frames, see the "Enabling Jumbo Frame Support (Supported Models)" section on page 15-37. 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. To use jumbo frames, set the value higher, for example, to 9000 bytes.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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 Click the **Advanced** tab.

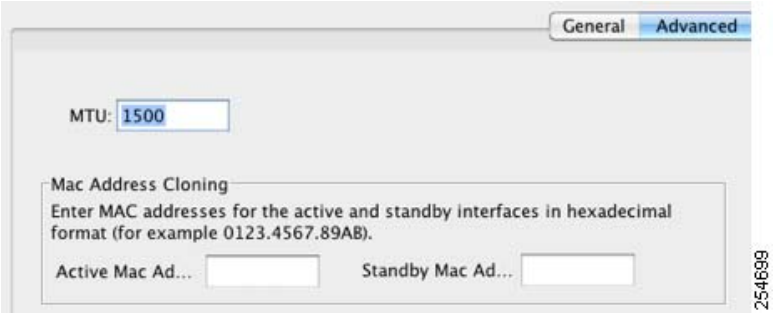

Step 4 To set the MTU or to enable jumbo frame support (supported models only), enter the value in the MTU field, between 300 and 65,535 bytes.

The default is 1500 bytes.

Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.

- For models that support jumbo frames 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 models that support jumbo frames 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 (Supported Models)" section on page 15-37.

Note Enabling or disabling jumbo frame support requires you to reload the ASA.

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. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

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 18-16.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the "IPv6 Addresses" section on page A-5.

This section includes the following topics:

- Information About IPv6, page 18-17

- Configuring a Global IPv6 Address, page 18-18
- Configuring IPv6 Neighbor Discovery, page 18-19
- (Optional) Configuring the Link-Local Addresses Automatically, page 18-19
- (Optional) Configuring the Link-Local Addresses Manually, page 18-20

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 18-17
- Modified EUI-64 Interface IDs, page 18-17
- Unsupported Commands, page 18-18

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. This address needs to be configured for each bridge group, and not per-interface. You can also configure a global IPv6 address for the management 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 address 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.

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 ASA can enforce this requirement for hosts attached to the local link.

When this feature 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:

```
%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.

Unsupported Commands

The following IPv6 commands are not supported in transparent firewall mode, because they require router capabilities:

- **ipv6 address autoconfig**
- **ipv6 nd prefix**
- **ipv6 nd ra-interval**
- **ipv6 nd ra-lifetime**
- **ipv6 nd suppress-ra**

Configuring a Global IPv6 Address

To configure a global IPv6 address for a bridge group or management interface, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately.

Restrictions

The ASA does not support IPv6 anycast addresses.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 16, "Starting Interface Configuration (ASA 5505)."
  - ASASM—Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the "Configuring Multiple Contexts" 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.

Step 2 Choose a BVI or management interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **IPv6** tab.

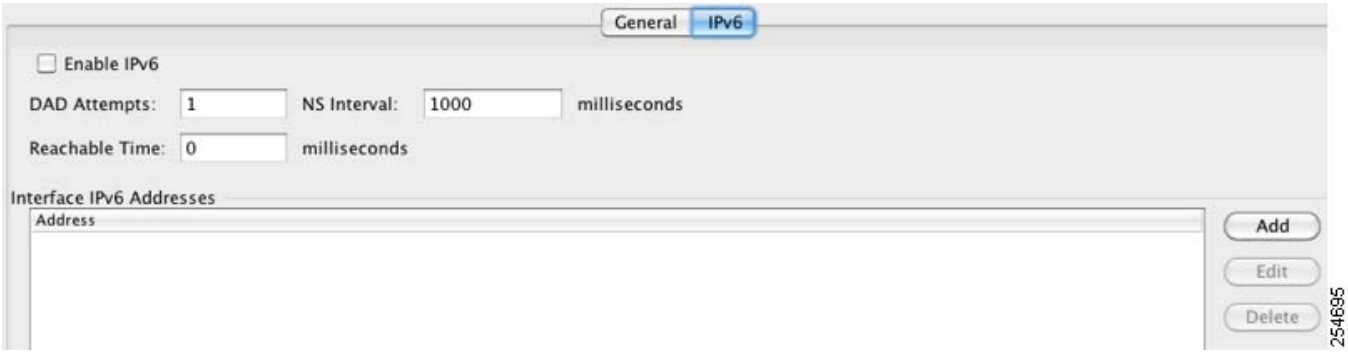

Step 4 Check the **Enable IPv6** check box.

Step 5 (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.

See the "Modified EUI-64 Interface IDs" section on page 18-17 for more information.

Step 6 (Optional) In the top area, customize the IPv6 configuration by referring to Chapter 37, "Configuring IPv6 Neighbor Discovery."

Step 7 To configure the 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 A-5 for more information about IPv6 addressing.

c. Click **OK**.

Step 8 Click **OK**.

You return to the Configuration > Device Setup > Interfaces pane.

Configuring IPv6 Neighbor Discovery

See Chapter 37, "Configuring IPv6 Neighbor Discovery," to configure IPv6 neighbor discovery.

(Optional) Configuring the Link-Local Addresses Automatically

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required for the interface ID.)

To manually assign the link-local address (not recommended), see the "(Optional) Configuring the Link-Local Addresses Manually" section on page 18-20.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 18-18.

To automatically configure the link-local addresses for a management interface or bridge group member interfaces, perform the following steps:

Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

Step 2 Select a BVI or management interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **IPv6** tab.

Step 4 In the IPv6 configuration area, check **Enable IPv6**.

This option enables IPv6 and automatically generates the link-local addresses for member interfaces using the Modified EUI-64 interface ID based on the interface MAC address.

Step 5 Click **OK**.

(Optional) Configuring the Link-Local Addresses Manually

If you do not want to configure a global address, and only need to configure a link-local address on the physical interfaces or subinterfaces, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

To automatically assign the link-local address (recommended), see the "(Optional) Configuring the Link-Local Addresses Automatically" section on page 18-19.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 18-18.

To assign a link-local address to a physical interface or subinterface, including the management 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 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 A-5 for more information about IPv6 addressing.

Step 5 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.

Information About Inter-Interface Communication

Allowing interfaces on the same security level to communicate with each other is useful if 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.

Detailed Steps

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**.

Turning Off and Turning On Interfaces

This section describes how to turn off and on an interface.

All interfaces are enabled by default. In multiple context mode, if you disable or reenable the interface within a context, only that context interface is affected. But if you disable or reenable the interface in the system execution space, then you affect that interface for all contexts.

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 VLAN interface that you want to configure, and click **Edit**.

The Edit Interface dialog box appears.

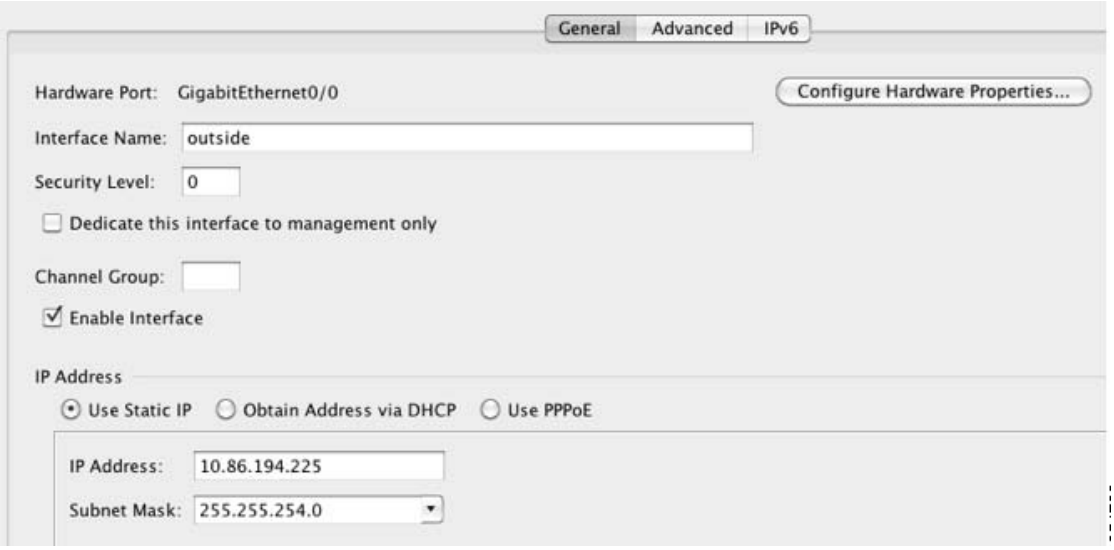

Step 3 To enable or disable the interface, check or uncheck the **Enable Interface** check box.

Monitoring Interfaces

This section includes the following topics:

- ARP Table, page 18-22
- DHCP, page 18-23
- MAC Address Table, page 18-25
- Dynamic ACLs, page 18-25
- Interface Graphs, page 18-26
- PPPoE Client, page 18-28
- Interface Connection, page 18-28

ARP Table

The Monitoring > Interfaces > 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.

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 ASA and updates Last Updated date and time.
- Last Updated—*Display only.* Shows the date and time the display was updated.

DHCP

The ASA lets you monitor DHCP status, including the addresses assigned to clients, the lease information for a ASA interface, and DHCP statistics.

DHCP Server Table

The Monitoring > Interfaces > DHCP > 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 ASA.
- Last Updated—Shows when the data in the table was last updated.

DHCP Client Lease Information

If you obtain the ASA interface IP address from a DHCP server, the Monitoring > Interfaces > DHCP > DHCP Server Table > DHCP Client Lease Information pane shows information about the DHCP lease.

Fields

- Select an interface—Lists the ASA 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 ASA begins the process of acquiring a lease. This state is also shown when a lease ends or when a lease negotiation fails.

    Selecting—The ASA is waiting to receive DHCPOFFER messages from one or more DHCP servers, so it can choose one.

    Requesting—The ASA is waiting to hear back from the server to which it sent its request.

    Purging—The ASA is removing the lease because of an error.

Bound—The ASA has a valid lease and is operating normally.

Renewing—The ASA is trying to renew the lease. It regularly sends DHCPREQUEST messages to the current DHCP server, and waits for a reply.

Rebinding—The ASA 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 ASA started the process to remove the lease.

Releasing—The ASA 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 ASA attempts to rebind to a DHCP server. Rebinding occurs if the ASA cannot communicate with the original DHCP server, and 87.5 percent of the lease time has expired. The ASA 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 ASA is attempting to establish a lease, this field shows the number of times the ASA tried sending a DHCP message. For example, if the ASA is in the Selecting state, this value shows the number of times the ASA sent discover messages. If the ASA is in the Requesting state, this value shows the number of times the ASA 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.

DHCP Statistics

The Monitoring > Interfaces > DHCP > 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 ASA.
- Total Messages Sent—Shows the total number of messages sent by the ASA.
- 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.

MAC Address Table

The Monitoring > Interfaces > MAC Address Table pane shows the static and dynamic MAC address entries. See the "MAC Address Table" section on page 18-25 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 the "MAC Address Table" section on page 18-25.
- Refresh—Refreshes the table with current information from the ASA.

Dynamic ACLs

The Monitoring > Interfaces > 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 ASA. 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.

Interface Graphs

The Monitoring > Interfaces > Interface Graphs pane lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA 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 ASA was incapable of handing received data to a hardware buffer because the input rate exceeded the ASA capability to handle the data.

    Underruns—The number of times that the transmitter ran faster than the ASA 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 ASA 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 ASA is partly finished sending the packet. The ASA 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.

- 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.

Graph/Table

The Monitoring > Interfaces > Interface Graphs > Graph/Table 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 (see the "Enabling History Metrics" section on page 4-34), 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 (see the "Enabling History Metrics" section on page 4-34). 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.

PPPoE Client

The Monitoring > Interfaces > PPPoE Client > 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 ASA for display.

Interface Connection

The Monitoring > Interfaces > *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 18-29
- Monitoring Statistics for, page 18-29

Track Status for

The Monitoring > Interfaces > interface connection > Track Status for pane displays information about 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.

Monitoring Statistics for

The Monitoring > Interfaces > interface connection > Monitoring Statistics 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.

Feature History for Interfaces in Transparent Mode

Table 18-1 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 18-1 Feature History for Interfaces in Transparent Mode

Feature Name	Platform 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 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 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 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.
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. We modified the following screen: Configuration > Device Setup > Interfaces > Switch Ports > Edit Switch Port.

Table 18-1 Feature History for Interfaces in Transparent Mode (continued)

Feature Name	Platform Releases	Feature Information
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames. 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. We modified the following screen: 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.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
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. We modified the following screens: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Bridge groups for transparent mode	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We modified or introduced the following screens: Configuration > Device Setup > Interfaces Configuration > Device Setup > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interfaces > Add/Edit Interface

Completing Interface Configuration (Transparent Mode, 8.3 and Earlier)

This chapter includes tasks to complete the interface configuration for all models in transparent firewall mode.

For Version 8.4 and later, see Chapter 19, "Completing Interface Configuration (Transparent Mode, 8.3 and Earlier)."

This chapter includes the following sections:

- Information About Completing Interface Configuration in Transparent Mode (8.3 and Earlier), page 19-1
- Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 19-3
- Guidelines and Limitations, page 19-3
- Default Settings, page 19-4
- Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier), page 19-4
- Completing Interface Configuration in Transparent Mode (8.3 and Earlier), page 19-9
- Monitoring Interfaces, page 19-17
- Feature History for Interfaces in Transparent Mode, page 19-18

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.

Information About Completing Interface Configuration in Transparent Mode (8.3 and Earlier)

This section includes the following topics:

- Information About the Global Management IP Address, page 19-2
- Security Levels, page 19-2

Information About the Global Management IP Address

A transparent firewall does not participate in IP routing. The only IP configuration required for the ASA is to set the management IP address. This address is required because the ASA uses this address as the source address for traffic originating on the ASA, 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 interface. This IP address can be on a separate subnet from the main management IP address.

Although you do not configure IPv4 or global IPv6 addresses for other interfaces, you still need to configure the security level and interface name according to the "Configuring General Interface Parameters" section on page 19-10.

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 19-17 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 19-17), 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 ASA.

- 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.

- **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.

Licensing Requirements for Completing Interface Configuration in Transparent Mode

Model	License Requirement
All models	Base License.

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 Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)." Then, configure the logical interface parameters in the context execution space according to this chapter.
- You can only configure context interfaces that you already assigned to the context in the system configuration.

Firewall Mode Guidelines

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

  Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire device. The ASA uses this IP address as the source address for packets originating on the ASA, such as system messages or AAA communications. In addition to the global management address, you can optionally configure a management interface; see the "Management Interface" section on page 15-2 for more information.

  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). The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported. See the "Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier)" section on page 19-4 for more information about management IP subnets.
- 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 ASA, you need to configure a global IPv6 address.
- 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.

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See the "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links. In multiple context mode, failover interfaces are configured in the system configuration.

IPv6 Guidelines

- Supports IPv6.
- No support for IPv6 anycast addresses in transparent mode.

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 3-17.

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 ASA 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.

Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier)

This section describes how to configure the management IP address for transparent firewall mode, and includes the following topics:

- Configuring the IPv4 Address, page 19-4
- Configuring the IPv6 Address, page 19-5

Configuring the IPv4 Address

This section tells how to configure the IPv4 address.

Detailed 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.

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

This section describes how to configure the global address or the link-local address, and includes the following topics:

- Information About IPv6, page 19-5
- Configuring the Global Address, page 19-7
- Configuring the Link-Local Addresses Automatically, page 19-7
- Configuring the Link-Local Address on an Interface Manually, page 19-8
- Configuring DAD Settings, page 19-8

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 19-5
- Duplicate Address Detection, page 19-6
- Modified EUI-64 Interface IDs, page 19-6
- Unsupported Commands, page 19-6

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. This address needs to be configured per device or context, and not per-interface. You can also configure a global IPv6 address for the management 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 address 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.

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:

```
%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 ASA uses neighbor solicitation messages to perform duplicate address detection. By default, the number of times an interface performs duplicate address detection is 1.

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 ASA can enforce this requirement for hosts attached to the local link.

When this feature 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:

```
%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.

Unsupported Commands

The following IPv6 commands are not supported in transparent firewall mode, because they require router capabilities:

- **ipv6 address autoconfig**
- **ipv6 nd prefix**
- **ipv6 nd ra-interval**

- **ipv6 nd ra-lifetime**
- **ipv6 nd suppress-ra**

The **ipv6 local pool** VPN command is not supported, because transparent mode does not support VPN.

Configuring the Global Address

To set the management IPv6 address, perform the following steps.

Detailed 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 A-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 A-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).

Detailed 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 19-8.

Step 3 Click **Apply**.

Configuring the Link-Local Address on an Interface Manually

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.

Detailed 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 "Modified EUI-64 Interface IDs" section on page 19-6 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 A-5 for more information about IPv6 addressing.

Step 6 Click **OK**.

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 19-7.

Detailed 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**.

Completing Interface Configuration in Transparent Mode (8.3 and Earlier)

This section includes tasks to complete the interface configuration for all models in transparent mode.

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:

- Task Flow for Completing Interface Configuration, page 19-9
- Configuring General Interface Parameters, page 19-10
- Configuring a Management Interface (ASA 5510 and Higher), page 19-11
- Configuring the MAC Address and MTU, page 19-15
- Allowing Same Security Level Communication, page 19-17

Task Flow for Completing Interface Configuration

Step 1 Complete the procedures in the "Starting Interface Configuration (ASA 5510 and Higher)" section on page 15-12 or the "Starting ASA 5505 Interface Configuration" section on page 16-6.

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 and security level. See the "Configuring General Interface Parameters" section on page 19-10.

Step 4 (Optional) Configure a management interface. See the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 19-11.

Step 5 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 19-15.

Step 6 (Optional) Allow same security level communication, either by allowing communication between two interfaces or by allowing traffic to enter and exit the same interface. See the "Allowing Same Security Level Communication" section on page 19-17.

Configuring General Interface Parameters

This procedure describes how to set the name, security level, and bridge group for each transparent interface.

To configure a separate management interface, see the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 19-11.

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

- You can configure up to two interfaces per context.
- For the ASA 5550 ASA, 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 19-2.
- 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 "Configuring Active/Standby Failover" section on page 12-7 or the "Configuring Active/Active Failover" section on page 13-8 to configure the failover and state links.

Prerequisites

- Complete the procedures in Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 16, "Starting Interface Configuration (ASA 5505)."
- 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.

In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for an interface, 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 19-2 for more information.

Step 5 If the interface is not already enabled, check the **Enable Interface** check box.

Step 6 (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, single mode) For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Step 7 Click **OK**.

What to Do Next

- (Optional) Configure a management interface. See the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 19-11.
- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 19-15.

Configuring a Management Interface (ASA 5510 and Higher)

You can configure one management interface separate from the network interfaces in single mode or per context. You can use the Management *slot*/*port* interface (either the physical interface or a subinterface) as a separate management interface. You cannot use any other interface type as a management interface. For more information, see the "Management Interface" section on page 15-2.

This section includes the following topics:

- Configuring General Parameters and the IPv4 Address, page 19-11
- Configuring a Global IPv6 Address and Other Options, page 19-13

Configuring General Parameters and the IPv4 Address

This section describes how to configure the name, security level, and IPv4 address for a management interface.

Prerequisites

- Complete the procedures in Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 16, "Starting Interface Configuration (ASA 5505)."
- 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.

In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a Management interface or subinterface 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 19-2 for more information.

Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

Step 5 If the interface is not already enabled, check the **Enable Interface** check box.

Step 6 To set the IP address, use one of the following options.

Note For use with failover, you must set the IP address and standby address manually; DHCP is not supported. Set the standby IP addresses on the Configuration > Device Management > High Availability > Failover > Interfaces tab.

- 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.

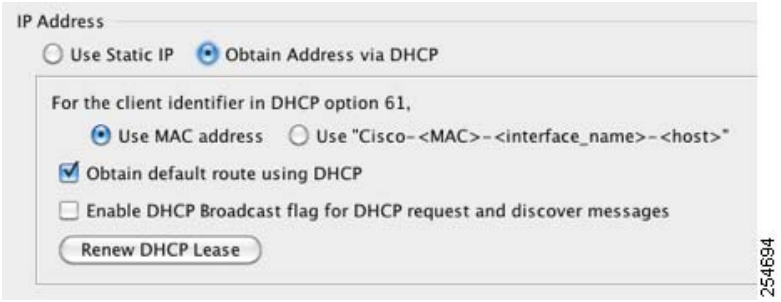

a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

   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. To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.

c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.

d. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

   The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

e. (Optional) To renew the lease, click **Renew DHCP Lease**.

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.

Note (ASA 5510 and higher, single mode) For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 15-23.

Step 8 Click **OK**.

Configuring a Global IPv6 Address and Other Options

To configure a global IPv6 address and other options for the management interface, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately.

Restrictions

The ASA does not support IPv6 anycast addresses.

Prerequisites

- Complete the procedures in Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 16, "Starting Interface Configuration (ASA 5505)."
- 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.

Step 2 Choose a management interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **IPv6** tab.

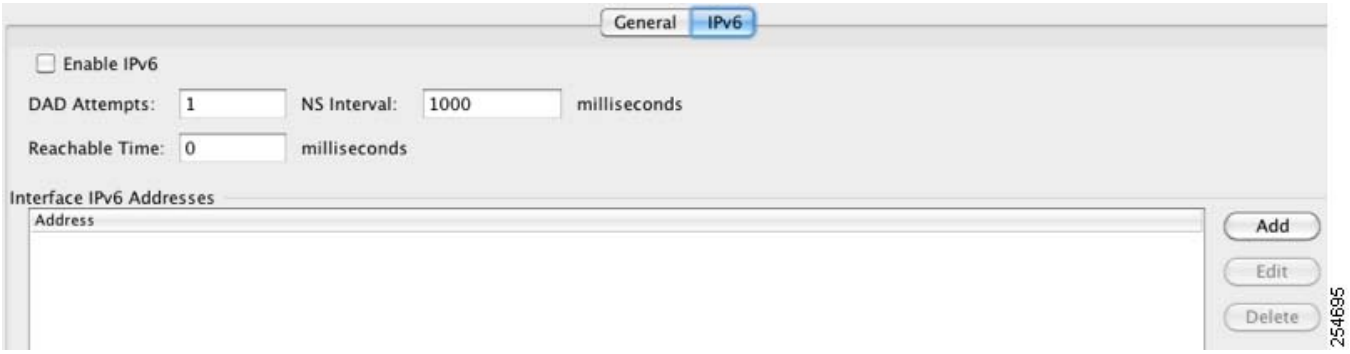

Step 4 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

See the "Modified EUI-64 Interface IDs" section on page 19-6 for more information.

Step 5 To configure the 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 A-5 for more information about IPv6 addressing.

c. 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.

Step 7 Click **OK**.

You return to the Configuration > Device Setup > Interfaces pane.

What to Do Next

(Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address and MTU" section on page 19-15.

Configuring the MAC Address and MTU

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.

For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. In multiple context mode, you can automatically assign unique MAC addresses to interfaces, including an EtherChannel port interface. We recommend manually, or in multiple context mode, automatically configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

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 ASA 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 ASA Classifies Packets" section on page 9-3 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 to Context Interfaces" section on page 9-24 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.

Prerequisites

- Complete the procedures in Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 16, "Starting Interface Configuration (ASA 5505)."
- 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.

BVIs appear in the table alongside physical interfaces, subinterfaces, redundant interfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a physical interface, subinterface, redundant interface, or EtherChannel port-interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **Advanced** tab.

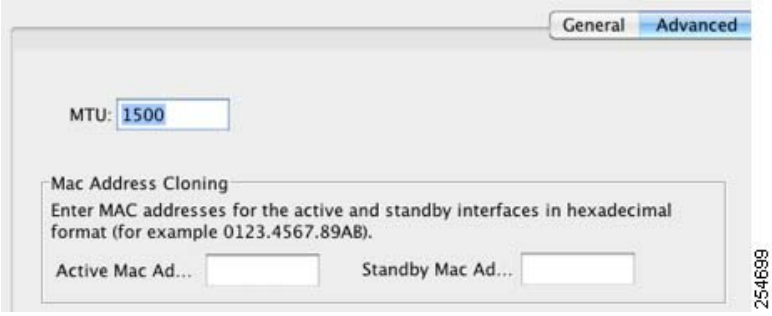

Step 4 To set the MTU or to enable jumbo frame support (ASA 5580 only), enter the value in the MTU field, between 300 and 65,535 bytes.

The default is 1500 bytes.

Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.

- 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 (Supported Models)" section on page 15-37.

Note Enabling or disabling jumbo frame support requires you to reboot the ASA.

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. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

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.

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.

Information About Inter-Interface Communication

Allowing interfaces on the same security level to communicate with each other is useful if 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.

Detailed Steps

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**.

Monitoring Interfaces

For information about monitoring screens, see the "Monitoring Interfaces" section on page 18-22.

Feature History for Interfaces in Transparent Mode

Table 19-1 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 19-1 Feature History for Interfaces in Transparent Mode

Feature Name	Platform 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 ASA 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 ASA (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 ASA (from 100 to 150), the ASA 5550 ASA (from 200 to 250).
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 ASA 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.
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. We modified the following screen: Configuration > Device Setup > Interfaces > Switch Ports > Edit Switch Port.

Table 19-1 Feature History for Interfaces in Transparent Mode (continued)

Feature Name	Platform Releases	Feature Information
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames. 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. We modified the following screen: 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.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
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. We modified the following screens: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Bridge groups for transparent mode (see Chapter 18, "Completing Interface Configuration (Transparent Mode, 8.4 and Later).")	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We modified or introduced the following screens: Configuration > Device Setup > Interfaces Configuration > Device Setup > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interfaces > Add/Edit Interface

PART 3

Configuring Basic Settings

Configuring Basic Settings

This chapter describes how to configure basic settings on the ASA that are typically required for a functioning configuration and includes the following sections:

- Configuring the Hostname, Domain Name, and Passwords, page 21-1
- Setting the Date and Time, page 21-2
- Configuring the Master Passphrase, page 21-4
- Configuring the DNS Server, page 21-8
- Changing the Heap Memory Size, page 21-9
- Monitoring DNS Cache, page 21-9

Configuring the Hostname, Domain Name, and Passwords

This section includes the following topic:

- Setting the Hostname, Domain Name, and the enable and Telnet Passwords, page 21-1

Setting the Hostname, Domain Name, and the enable and Telnet Passwords

To set the hostname, domain name, and the enable and Telnet passwords, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Device Name/Password**.

Step 2 Enter the hostname. The default hostname depends on your platform.

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 syslog 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.

Step 3 Enter the domain name. The default domain name is default.domain.invalid.

The ASA 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 ASA qualifies the name to “jupiter.example.com.”

Note In multiple context mode, the Enable Password area only appears in contexts; it does not appear in the system execution space.

Step 4 Change the privileged mode (enable) password.

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 a separate password, and this enable password is not used. In addition, you can configure authentication for HTTP/ASDM access.)

Step 5 Enter the old password.

Step 6 Enter the new password.

Step 7 Confirm the new password.

Note In multiple context mode, the Telnet Password area only appears in contexts; it does not appear in the system execution space.

Step 8 Change the login password to access the platform console.

The Telnet password sets the login password. By default, it is "cisco." This password applies to both Telnet and SSH access. The login password lets you access EXEC mode if you connect to the ASA using a Telnet or SSH session. (If you configure user authentication for Telnet or SSH access, then each user has a separate password, and this login password is not used.)

Step 9 Enter the old password.

Step 10 Enter the new password.

Step 11 Confirm the new password.

Step 12 Click **Apply** to save your changes.

Setting the Date and Time

Note Do not set the date and time for the ASASM; it receives these settings from the host switch.

This section includes the following topics:

- Setting the Date and Time Using an NTP Server, page 21-3
- Setting the Date and Time Manually, page 21-4

Setting the Date and Time Using an NTP Server

To obtain the date and time from an NTP server, choose Configuration > Device Setup > System Time > NTP:

Detailed Steps

Use the NTP pane to define NTP servers for setting the time dynamically on the ASA. The time appears in the status bar at the bottom of the main ASDM window. Time derived from an NTP server overrides any time set manually in the Clock 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 ASA chooses the server with the lowest stratum—a measure of how reliable the data is.

Adding or Editing the NTP Server Configuration

To add or edit an NTP server, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > System Time > NTP**.

Step 2 Click **Add** to display the Add NTP Server Configuration dialog box.

Step 3 Enter the NTP server IP address.

Step 4 Check the **Preferred** check box to set this server as a preferred server. NTP uses an algorithm to determine which server is the most accurate and synchronizes to it. 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 ASA uses the more accurate one.

Step 5 Choose the interface from the drop-down list. This setting specifies the outgoing interface for NTP packets. If the interface is blank, then the ASA uses the default admin context interface according to the routing table. To change the admin context (and the available interfaces), choose None (the default interface) for stability.

Step 6 Choose the key number from the drop-down list. This setting specifies the key ID for this authentication key, which enables you to use MD5 authentication to communicate with the NTP server. The NTP server packets must also use this key ID. If you have previously configured a key ID for another server, you can select it from the list; otherwise, enter a number between 1 and 4294967295.

Step 7 Check the **Trusted** check box to set this authentication key as a trusted key, which is required for authentication to succeed.

Step 8 Enter the key value to set the authentication key, which is a string that can be up to 32 characters long.

Step 9 Reenter the key value to make sure that you enter it correctly twice.

Step 10 Click **OK**.

Step 11 Check the **Enable NTP authentication** check box to turn on NTP authentication.

Step 12 Click **Apply** to save your changes.

Setting the Date and Time Manually

The time is based on a 24-hour clock and 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, choose **Configuration > Device Setup > System Time > NTP**; time derived from an NTP server overrides any time set manually in the Clock pane.

To manually set the date and time for the ASA, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > System Time > Clock**.

Step 2 Choose the time zone from the drop-down list. This setting specifies 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 ASA may drop the connection to intelligent SSMs.

Step 3 Click the Date drop-down list to display a calendar. Then find the correct date using the following methods:

- Click the name of the month to display a list of months, then click the desired month. The calendar updates to that month.
- Click the year to change the year. Use the up and down arrows to scroll through the years, or enter a year in the entry field.
- Click the arrows to the right and left of the month and year to scroll the calendar forward and backward one month at a time.
- Click a day on the calendar to set the date.

Step 4 Enter the time manually in hours, minutes, and seconds.

Step 5 Click **Update Display Time** to update the time shown in the bottom right corner of the ASDM pane. The current time updates automatically every ten seconds.

Configuring the Master Passphrase

This section includes the following topics:

- Information About the Master Passphrase, page 21-5
- Licensing Requirements for the Master Passphrase, page 21-5
- Guidelines and Limitations, page 21-5
- Adding or Changing the Master Passphrase, page 21-5
- Disabling the Master Passphrase, page 21-6
- Recovering the Master Passphrase, page 21-7
- Feature History for the Master Passphrase, page 21-7

Information About the Master Passphrase

The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. Features that use the master passphrase include the following:

- OSPF
- EIGRP
- VPN load balancing
- VPN (remote access and site-to-site)
- Failover
- AAA servers
- Logging
- Shared licenses

Licensing Requirements for the Master Passphrase

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.

Failover Guidelines

If failover is enabled but no failover shared key is set, an error message appears if you change the master passphrase, informing you that you must enter a failover shared key to protect the master passphrase changes from being sent as plain text.

Choose **Configuration > Device Management > High Availability > Failover**, enter any character in the Shared Key field or 32 hexadecimal numbers (0-9A-Fa-f) if a failover hexadecimal key is selected, except a backspace. Then click **Apply.**

Adding or Changing the Master Passphrase

To add or change the master passphrase, perform the following steps:

Step 1 In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.

In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase.**

Step 2 Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning message appears when you click **Apply**. You can click **OK** or **Cancel** to continue.

If you later disable password encryption, all existing encrypted passwords are left unchanged, and as long as the master passphrase exists, the encrypted passwords will be decrypted as required by the application.

Step 3 Check the **Change the encryption master passphrase** check box to enable you to enter and confirm your new master passphrases. By default, they are disabled.

Your new master passphrase must be between 8 and 128 characters long.

If you are changing an existing passphrase, you must enter the old passphrase before you can enter a new one.

To delete the master passphrase, leave the New and Confirm master passphrase fields blank.

Step 4 Click **Apply**.

When you click **Apply**, warning messages appear under the following conditions:

- The Change the encryption master passphrase field is enabled, and the New master passphrase field is empty. The **no key configuration-key password-encrypt** command is then sent to the device.
- The old master passphrase does not match the hash value in the **show password encryption** command output.
- You use non-portable characters, particularly those with the high-order bit set in an 8-bit representation.
- A master passphrase and failover are in effect, then an error message appears if an attempt to remove the failover shared key occurs.
- Encryption is disabled, but a new or replacement master passphrase is supplied. Click **OK** or **Cancel** to continue.
- If the master passphrase is changed in multiple context mode.
- If Active/Active failover is configured and the master passphrase is changed.
- If any running configurations are configured so that their configurations cannot be saved to their server, such as with context configuration URLs that use HTTP or HTTPS, and the master passphrase is changed.

Disabling the Master Passphrase

Disabling the master passphrase reverts encrypted passwords into plain text passwords. Removing the passphrase might be useful if you downgrade to a previous software version that does not support encrypted passwords.

You must know the current master passphrase to disable it. If you do not know the passphrase, see the "Recovering the Master Passphrase" section on page 21-7.

To disable the master passphrase, perform the following steps:

Step 1 In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.

In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.

Step 2 Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning statement appears when you click **Apply**. Click **OK** or **Cancel** to continue.

Step 3 Check the **Change the encryption master passphrase** check box.

Step 4 Enter the old master passphrase in the Old master passphrase field. You must provide the old master passphrase to disable it.

Step 5 Leave the New master passphrase and the Confirm master passphrase fields empty.

Step 6 Click **Apply**.

Recovering the Master Passphrase

You cannot recover the master passphrase. If the master passphrase is lost or unknown, you can remove it.

To remove the master passphrase, perform the following steps:

Feature History for the Master Passphrase

Table 21-1 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 21-1 Feature History for the Master Passphrase

Feature Name	Platform Releases	Feature Information
Master Passphrase	8.3(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management > Advanced > Master Passphrase. Configuration > Device Management > Device Administration > Master Passphrase.

Configuring the DNS Server

Some ASA features require use of a DNS server to access external servers by domain name; for example, the Botnet Traffic Filter feature requires a DNS server to access the dynamic database server and to resolve entries in the static database. Other features, such as the **ping** or **traceroute** command, let you enter a name that you want to ping or traceroute, and the ASA can resolve the name by communicating with a DNS server. Many SSL VPN and certificate commands also support names.

Note The ASA has limited support for using the DNS server, depending on the feature. For these feature, to resolve the server name to an IP address, you must enter the IP address manually by adding the server name in the Configuration > Firewall > Objects > Network Object/Groups pane.

For information about dynamic DNS, see the "Configuring Dynamic DNS" section on page 22-2.

Prerequisites

Make sure that you configure the appropriate routing for any interface on which you enable DNS domain lookup so you can reach the DNS server. See the "Information About Routing" section on page 30-1 for more information about routing.

Step 1 In the ASDM main application window, choose **Configuration > Device Management > DNS > DNS Client**.

Step 2 In the DNS Setup area, choose one of the following options:

- Configure one DNS server group.
- Configure multiple DNS server groups.

Step 3 Click **Add** to display the Add DNS Server Group dialog box.

Step 4 Specify up to six addresses to which DNS requests can be forwarded. The ASA tries each DNS server in order until it receives a response.

Note You must first enable DNS on at least one interface before you can add a DNS server. The DNS Lookup area shows the DNS status of an interface. A False setting indicates that DNS is disabled. A True setting indicates that DNS is enabled.

Step 5 Enter the name of each configured DNS server group.

Step 6 Enter the IP addresses of the configured servers, and click **Add** to include them in the server group. To remove a configured server from the group, click **Delete**.

Step 7 To change the sequence of the configured servers, click **Move Up** or **Move Down**.

Step 8 In the Other Settings area, enter 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 ASA retries the list of servers, the timeout time doubles.

Step 9 Enter the number of seconds to wait before trying the next DNS server in the group.

Step 10 Enter a valid DNS domain name for the group of configured servers.

Step 11 Click **OK** to close the Add DNS Server Group dialog box.

The new DNS server settings appear.

Step 12 To change these settings, click **Edit** to display the Edit DNS Server Group dialog box.

Step 13 Make your desired changes, then click **OK** to close the Edit DNS Server Group dialog box.

The revised DNS server settings appear.

Step 14 To enable a DNS server group to receive DNS requests, click **Set Active**.

Step 15 In the DNS Guard area, to enforce one DNS response per query, check the **Enable DNS Guard on all interfaces** check box. If DNS inspection is enabled, this setting is ignored on the selected interface.

Step 16 Click **Apply** to save your changes, or click **Reset** to discard those changes and enter new ones.

Changing the Heap Memory Size

Along with using troubleshooting information in this guide, see the *ASDM Troubleshooting* document at the following URL:

http://www.cisco.com/en/US/products/ps6121/products_tech_note09186a0080aaeff5.shtml

ASDM supports a maximum configuration size of 512 KB. If you exceed this amount you may experience performance issues. For example, when you load the configuration, the status dialog box shows the percentage of the configuration that is complete, yet with large configurations it stops incrementing and appears to suspend operation, even though ASDM might still be processing the configuration. If this situation occurs, we recommend that you consider increasing the ASDM system heap memory.

To increase the ASDM heap memory size, modify the launcher shortcut by performing the following procedure:

Step 1 Right-click the shortcut for the ASDM-IDM Launcher, and choose **Properties**.

Step 2 Click the **Shortcut** tab.

Step 3 In the Target field, change the argument prefixed with "-Xmx" to specify your desired heap size. For example, change it to -Xmx768m for 768 MB or -Xmx1g for 1 GB. For more information about this parameter, see the Oracle document in the following location:
http://docs.oracle.com/javase/1.5.0/docs/tooldocs/windows/java.html

Monitoring DNS Cache

The ASA 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 with its corresponding hostname.

To monitor the DNS cache, see the following pane:

Path	Purpose
Tools > Command Line Interface Enter the **show** dns-hosts command, then press **Send.**	Show the DNS cache, which includes dynamically learned entries from a DNS server as well as manually entered name and IP addresses using the **name** command.

Configuring Dynamic DNS

This chapter describes how to configure DDNS update methods and includes the following topics:

- Information About DDNS, page 22-1
- Licensing Requirements for DDNS, page 22-2
- Guidelines and Limitations, page 22-2
- Configuring Dynamic DNS, page 22-2
- DDNS Monitoring, page 22-4
- Feature History for DDNS, page 22-4

Information About DDNS

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 predefined 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. DDNS provides the necessary dynamic update and synchronization of the name-to-address mapping and address-to-name mapping on the DNS server. To configure the DNS server for other uses, see the "Configuring the DNS Server" section on page 21-8. To configure DHCP, see the "Configuring a DHCP Server" section on page 23-3.

EDNS allows DNS requesters to advertise the size of their UDP packets and facilitates the transfer of packets larger than 512 octets. When a DNS server receives a request over UDP, it identifies the size of the UDP packet from the OPT resource record (RR) and scales its response to contain as many resource records as are allowed in the maximum UDP packet size specified by the requester. The size of the DNS packets can be up to 4096 bytes for BIND or 1280 bytes for the Windows 2003 DNS Server. Several additional **message-length maximum** commands are available:

- The existing global limit: **message-length maximum 512**
- A client or server specific limit: **message-length maximum client 4096**
- The dynamic value specified in the OPT RR field: **message-length maximum client auto**

If the three commands are present at the same time, the ASA enforces the minimum of the three specified values.

Licensing Requirements for DDNS

The following table shows the licensing requirements for DDNS:

Model	License Requirement
All models	Base License.

Guidelines and Limitations

Failover Guidelines

Supports Active/Active and Active/Standby failover.

Firewall Mode Guidelines

Supported in routed firewall mode.

Context Mode Guidelines

Supported in single and multiple context modes.

Supported in transparent mode for the DNS Client pane.

IPv6 Guidelines

Supports IPv6.

Configuring Dynamic DNS

Dynamic DNS provides address and domain name mapping so that hosts can find each other, even though their DHCP-assigned IP addresses change frequently. The DDNS name and address mapping are stored on the DHCP server in two resource records: the A RR includes 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 ASA supports the IETF method in this release.

The Dynamic DNS pane shows the configured DDNS update methods and the interfaces that have been configured for DDNS. By automatically recording the association between assigned addresses and hostnames at pre-defined intervals, DDNS allows frequently changing address-hostname associations to be updated regularly. Mobile hosts, for example, can then move freely on a network without user or administrator intervention.

To configure dynamic DNS client settings for updating the DNS server, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Device Management > DNS > Dynamic DNS**.

Step 2 Click **Add** to display the Add Dynamic DNS Update Method dialog box.

Step 3 Enter the name for the DDNS update method.

Step 4 Specify the update interval between DNS update attempts configured for the update method in days, hours, minutes, and seconds.

- Choose the number of days between update attempts from 0 to 364.
- Choose the number of hours (in whole numbers) between update attempts from 0 to 23.
- Choose the number of minutes (in whole numbers) between update attempts from 0 to 59.
- Choose the number of seconds (in whole numbers) between update attempts from 0 to 59.

These units are additive. That is, if you enter 0 days, 0 hours, 5 minutes and 15 seconds, the update method tries an update every 5 minutes and 15 seconds for as long as the method is active.

Step 5 To store server resource record updates that the DNS client updates, choose one of the following options:

- Both the A resource record and the PTR resource record.
- The A resource records only.

Step 6 Click **OK** to close the Add Dynamic DNS Update Method dialog box.

The new dynamic DNS client settings appear.

Step 7 To change these settings, click **Edit** to display the Edit Dynamic DDNS Update Method dialog box. When you edit an existing method, the Name field is *display-only* and shows the name of the selected method for editing

Step 8 Make your desired changes, and then click OK to close the Edit Dynamic DDNS Update Method dialog box.

The revised dynamic DNS client settings appear.

Step 9 To remove configured settings, choose the settings from the list, and then click **Delete**.

Step 10 To add DDNS settings for each interface configured for DDNS, click **Add** to display the Add Dynamic DNS Interface Settings dialog box.

Step 11 Choose the interface from the drop-down list.

Step 12 Choose the update method assigned to the interface from the drop-down list.

Step 13 Enter the hostname of the DDNS client.

Step 14 To store resource record updates, choose one of the following options:

- Default (PTR Records) to specify that the client request PTR record updating by the server.
- Both (PTR Records and A Records) to specify that the client request both the A and PTR DNS resource records by the server.
- 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 take effect.

Step 15 Click **OK** to close the Add Dynamic DNS Interface Settings dialog box.

The new dynamic DNS interface settings appear.

Step 16 To change these settings, click **Edit** to display the Edit Dynamic DNS Interface Settings dialog box.

Step 17 Make your desired changes, and then click **OK** to close the Edit Dynamic DNS Interface Settings dialog box.

The revised dynamic DNS interface settings appear.

Step 18 To remove configured settings, choose the settings from the list, and then click **Delete**.

Step 19 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

DDNS Monitoring

To monitor DDNS, perform the following steps:

Path	Purpose
Tools > Command Line Interface Enter the **showrunning-config ddns** command, then click **Send**.	Shows the current DDNS configuration.
Tools > Command Line Interface Enter the **show running-config dns server-group** command, then click **Send**.	Shows the current DNS server group status.

Feature History for DDNS

Table 22-1 lists each feature change and the platform release in which it was implemented.

ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 22-1 Feature History for DDNS

Feature Name	Releases	Feature Information
DDNS	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management> DNS > DNS Client. Configuration > Device Management > DNS > Dynamic DNS.

Configuring DHCP

This chapter describes how to configure the DHCP server and includes the following sections:

- Information About DHCP, page 23-1
- Licensing Requirements for DHCP, page 23-2
- Guidelines and Limitations, page 23-2
- Configuring the DHCP Relay Service, page 23-6
- Configuring a DHCP Server, page 23-3
- Additional References, page 23-8
- DHCP Monitoring, page 23-9
- Feature History for DHCP, page 23-10

Information About DHCP

The DHCP Relay Agent sends Dynamic Host Configuration Protocol (DHCP) messages between DHCP clients and servers on different IP networks. DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The ASA can provide a DHCP server or DHCP relay service to DHCP clients attached to ASA interfaces. The DHCP server provides network configuration parameters directly to DHCP clients. The DHCP relay service sends DHCP requests received on one interface to an external DHCP server located on a different interface.

A client locates a DHCP server to request the assignment of configuration information using a reserved, link-scoped multicast address, which indicates that the client and server should be attached to the same link. However, in some cases where ease of management, economy, or scalability is the concern, we recommend that you allow a DHCP client to send a message to a server that is not connected to the same link. The DHCP relay agent, which may reside on the client network, can relay messages between the client and server. The relay agent operation is transparent to the client.

DHCP for IPv6 (DHCPv6) specified in RFC 3315 enables IPv6 DHCP servers to send configuration parameters such as network addresses or prefixes and DNS server addresses to IPv6 nodes (that is, DHCP clients). DHCPv6 uses the following multicast addresses:

- All_DHCP_Relay_Agents_and_Servers (FF02::1:2) is a link-scoped multicast address used by a client to communicate with neighboring (that is, on-link) relay agents and servers. All DHCPv6 servers and relay agents are members of this multicast group.
- The DHCPv6 relay service and server listen for messages on UDP port 547. The ASA DHCPv6 relay agent listens on both UDP port 547 and the All_DHCP_Relay_Agents_and_Servers multicast address.

Licensing Requirements for DHCP

Table 23-1 shows the licensing requirements for DHCP.

Table 23-1 Licensing Requirements

Model	License Requirement
All models	Base License.

For the ASA 5505, the maximum number of DHCP client addresses varies depending on the license:

- If the limit is 10 hosts, the maximum available DHCP pool is 32 addresses.
- If the limit is 50 hosts, the maximum available DHCP pool is 128 addresses.
- If the number of hosts is unlimited, the maximum available DHCP pool is 256 addresses.

Note The ASA 5505 ships with a 10-user license.

Guidelines and Limitations

Use the following guidelines to configure the DHCP server:

- You can configure only one DHCP server on each interface of the ASA. Each interface can have its own pool of addresses to use. However the other DHCP settings, such as DNS servers, domain name, options, ping timeout, and WINS servers, are configured globally and used by the DHCP server on all interfaces.
- You cannot configure a DHCP client or DHCP relay service on an interface on which the server is enabled. Additionally, DHCP clients must be directly connected to the interface on which the server is enabled.
- The ASA does not support QIP DHCP servers for use with the DHCP proxy service.
- The relay agent cannot be enabled if the DHCP server is also enabled.
- The ASA DHCP server does not support BOOTP requests. In multiple context mode, you cannot enable the DHCP server or DHCP relay service on an interface that is used by more than one context.
- When it receives a DHCP request, the ASA sends a discovery message to the DHCP server. This message includes the IP address (within a subnetwork) that was configured with the **dhcp-network-scope** command in the group policy. If the server has an address pool that falls within that subnetwork, the server sends the offer message with the pool information to the IP address—not to the source IP address of the discovery message.

  For example, if the server has a pool in the range of 209.165.200.225 to 209.165.200.254, mask 255.255.255.0, and the IP address specified by the **dhcp-network-scope** command is 209.165.200.1, the server sends that pool in the offer message to the ASA.

Use the following guidelines to configure the DHCP relay service:

- DHCP clients must be directly connected to the ASA and cannot send requests through another relay agent or a router.

- For multiple context mode, you cannot enable DHCP relay service on an interface that is used by more than one context.
- The DHCP relay service is not available in transparent firewall mode. An ASA in transparent firewall mode only allows ARP traffic through; all other traffic requires an access list. To allow DHCP requests and replies through the ASA in transparent firewall mode, you must configure two 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.
- When the DHCP relay service is enabled and more than one DHCP relay server is defined, the ASA forwards client requests to each defined DHCP relay server. Replies from the servers are also forwarded to the client until the client DHCP relay binding is removed. The binding is removed when the ASA receives any of the following DHCP messages: ACK, NACK, or decline.
- You cannot enable DHCP relay service on an interface running as a DHCP proxy service. You must remove the VPN DHCP configuration first or an error message appears. This error occurs if both DHCP relay and DHCP proxy services are enabled. Make sure that either the DHCP relay or DHCP proxy service is enabled, but not both.

Firewall Mode Guidelines

Supported in routed firewall mode.

Not supported in transparent firewall mode.

Context Mode Guidelines

Supported in single and multiple context mode.

Failover Guidelines

Supports Active/Active and Active/Standby failover.

IPv6 Guidelines

Supports IPv6.

Configuring a DHCP Server

This section includes the following topics:

- Editing DHCP Server Settings, page 23-4
- Configuring Advanced DHCP Options, page 23-5

To configure an ASA interface as a DHCP server, perform the following steps:

Step 1 In the Global DHCP Options area, check the **Enable Auto-configuration from interface** check box to enable DHCP auto configuration and choose the interface from the drop-down list.

DHCP auto configuration enables the DHCP Server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. If information obtained through auto configuration is also specified manually in the Global DHCP Options area, the manually specified information takes precedence over the discovered information.

Step 2 To override the interface DHCP or PPPoE client WINS parameter with the VPN client parameter, check the **Allow VPN override** check box.

Step 3 In the DNS Server 1 field, enter the IP address of the primary DNS server for a DHCP client.

Step 4 In the DNS Server 2 field, enter the IP address of the alternate DNS server for a DHCP client.

Step 5 In the Domain Name field, enter the DNS domain name for DHCP clients (for example, example.com).

Step 6 In the Lease Length field, 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).

Step 7 In the Primary WINS Server field, enter the IP address of the primary WINS server for a DHCP client.

Step 8 In the Secondary WINS Server field, enter the IP address of the alternate WINS server for a DHCP client.

Step 9 To avoid address conflicts, the ASA sends two ICMP ping packets to an address before assigning that address to a DHCP client. In the Ping Timeout field, enter the amount of time, in milliseconds, that the ASA waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.

Step 10 To specify additional DHCP options and their parameters, click **Advanced** to display the Configuring Advanced DHCP Options dialog box. For more information, see the "Configuring Advanced DHCP Options" section on page 23-5.

Step 11 In the Dynamic DNS Settings for DHCP Server area, you configure the DDNS update settings for the DHCP server. Check the **Update DNS Clients** check box to specify that, in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:

- Check the **Update Both Records** check box to specify that the DHCP server should update both the A and PTR RRs.
- Check the **Override Client Settings** check box to specify that the DHCP server actions should override any update actions requested by the DHCP client.

Step 12 To modify DHCP server settings, click **Edit** to display the Edit DHCP Server dialog box. Alternatively, you may double-click the row for a particular interface to open the Edit DHCP Server dialog box for that interface. For more information, see the "Editing DHCP Server Settings" section on page 23-4.

Step 13 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Editing DHCP Server Settings

To enable DHCP, specify the DHCP address pool, and modify other DHCP server parameters for the selected interface, perform the following steps:

Step 1 The selected interface ID appears as *display-only*. To enable the DHCP server on the selected interface, check the **Enable DHCP Server** check box. To disable DHCP on the selected interface, uncheck this check box. Disabling the DHCP server on the selected interface does not clear the specified DHCP address pool.

Step 2 In the DHCP Address Pool field, enter the range of IP addresses from lowest to highest that is used by the DHCP server. The range of IP addresses must be on the same subnet as the selected interface and cannot include the IP address of the interface itself.

Step 3 Check the **Update DNS Clients** check box to specify that, in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:

- To specify that the DHCP server should update both the A and PTR RRs, check the **Update Both Records** check box.

- To specify that DHCP server actions should override any update actions requested by the DHCP client, check the **Override Client Settings** check box

Step 4 To enable DHCP on the interface, check the **Enable DHCP Server** check box. The DHCP Enabled column displays "Yes" if DHCP is enabled, or "No" if DHCP is disabled on the interface.

Step 5 In the DNS Address Pool field, enter the revised range of IP addresses that have been assigned to the DHCP address pool.

Step 6 In the Optional Parameters area, modify the following settings:

a. The DNS servers (1 and 2) configured for the interface.

b. The WINS servers (primary and secondary) configured for the interface.

c. The domain name of the interface.

d. The time in milliseconds that the ASA will wait for an ICMP ping response on the interface.

e. The duration of time that the DHPC server configured on the interface allows DHCP clients to use an assigned IP address.

f. The interface on a DHCP client that provides DNS, WINS, and domain name information for automatic configuration.

Step 7 To configure more DHCP options, click **Advanced** to display the Advanced DHCP Options dialog box. For more information, see the "Configuring Advanced DHCP Options" section on page 23-5.

Step 8 Click **OK** to close the Edit DHCP Server dialog box.

Step 9 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Configuring Advanced DHCP Options

You can use advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients. You can also use the DHCP automatic configuration setting to obtain these values or define them manually. When you use more than one method to define this information, it is passed to DHCP clients in the following sequence:

1. Manually configured settings.
2. Advanced DHCP options settings.
3. DHCP automatic configuration settings.

For example, you can manually define the domain name that you want the DHCP clients to receive and then enable DHCP automatic configuration. Although DHCP automatic configuration discovers the domain together with the DNS and WINS servers, the manually defined domain name is passed to DHCP clients with the discovered DNS and WINS server names, because the domain name discovered by the DHCP automatic configuration process is superseded by the manually defined domain name.

To configure advanced DHCP options, perform the following steps:

Step 1 In the Option to be Added area, define the following settings:

a. Choose the option code from the drop-down list. All DHCP options (options 1 through 255) are supported except 1, 12, 50–54, 58–59, 61, 67, and 82.

b. Choose the options 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 example, if you choose DHCP Option 2 (Time Offset), you can only enter 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.

c. In the Option Data area, specify the type of information that 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. Click **Add** to add the option to the DHCP option list. Click **Delete** to remove the option from the DHCP option list.

   – Click **IP Address** to indicate that an IP address is returned to the DHCP client. You can specify up to two IP addresses. IP Address 1 and IP Address 2 indicate an IP address in dotted-decimal notation.

Note The name of the associated IP address fields can change based on the DHCP option that you chose. For example, if you choose DHCP Option 3 (Router), the fields names change to Router 1 and Router 2.

   – Click **ASCII** to specify that an ASCII value is returned to the DHCP client. In the Data field, enter an ASCII character string. The string cannot include spaces.

Note The name of the associated Data field can change based on the DHCP option that you chose. For example, if you choose DHCP Option 14 (Merit Dump File), the associated Data field names change to File Name.

   – Click **Hex** to specify that a hexadecimal value is returned to the DHCP client. In the Data field, enter a hexadecimal string with an even number of digits and no spaces. You do not need to use a 0x prefix.

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.

Step 2 Click **OK** to close the Advanced DHCP Options dialog box.

Step 3 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Configuring the DHCP Relay Service

The DHCP Relay pane lets you configure the DHCP relay service on the ASA through ASDM.

This section includes the following topics:

- Editing DHCP Relay Agent Settings, page 23-7
- Adding or Editing Global DHCP Relay Server Settings, page 23-8

To configure the DHCP relay service, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Device Management > DHCP > DHCP Relay**.

Step 2 To configure a new external server to which DHCP requests may be relayed, click **Add** to display the Add Global DHCP Relay Server dialog box. You can define up to ten global DHCP relay servers for IPv4 and ten global DHCP relay servers for IPv6 on the ASA. If you already have defined ten DHCP relay global servers for IPv4 and ten DHCP relay global servers for IPv6, the Add button is dimmed.

Step 3 Enter the IP address of the DHCP server, or click the ellipses to display the Browse DHCP Server dialog box.

Step 4 Double-click a DHCP server from the list to add it to the DHCP Server field, then click **OK** to close the Browse DHCP Server dialog box.

The newly selected DHCP server appears in the DHCP Server field.

Step 5 Choose the interface to which the specified DHCP server is attached from the drop-down list, then click **OK** to close the Add Global DHCP Relay Server dialog box.

The newly added global DHCP relay server appears in the Global DHCP Relay Servers list.

Step 6 To change a selected DHCP relay server settings, click **Edit** to display the Edit Global DHCP Relay Server dialog box.

Step 7 Make your desired changes, then click **OK** to close the Edit Global DHCP Relay Server dialog box. To remove a selected DHCP relay server from the list, click **Delete**.

Step 8 The DHCP Relay Agent area shows the following information:

- The configured interfaces.
- The IPv4 DHCP Relay Enabled column and IPv6 DHCP Relay Enabled column indicate whether or not the selected DHCP relay agent is enabled on the interface. The check box is checked if the DHCP relay agent is enabled, and is unchecked if the DHCP relay agent is not enabled on the interface.
- The Set Route column indicates whether or not the selected DHCP relay agent is configured to modify the default router address in the information returned from the DHCP server. The check box is checked if the DHCP relay agent is configured to change the default router address to the interface address, and is unchecked if the DHCP relay agent does not modify the default router address.

Step 9 Enter the amount of time, in seconds, allowed for DHCP address handling in the IPv4 timeout or IPv6 timeout field. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

Step 10 Click **Apply** to save your settings, or click **Reset** to discard them and enter new ones.

Step 11 To change DHCP relay agent settings, see the "Editing DHCP Relay Agent Settings" section on page 23-7.

Editing DHCP Relay Agent Settings

To modify DHCP relay agent settings, perform the following steps:

Step 1 To enable the DHCP relay agent on the selected interface, check the **Enable DHCP Relay Agent** check box.

Step 2 To specify whether the DHCP relay agent is configured to modify the default router address in the information returned from the DHCP server, check the **Set Route** check box. The DHCP relay agent then substitutes the address of the selected interface for the default router address in the information returned from the DHCP server.

Step 3 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Step 4 To add or edit global DHCP relay server settings, see the "Adding or Editing Global DHCP Relay Server Settings" section on page 23-8.

Adding or Editing Global DHCP Relay Server Settings

To add or edit a global DHCP relay server to which DHCP requests are relayed, perform the following steps:

Step 1 In the Global DHCP Relay Servers area, click **Add** to display the Add Global Relay Servers dialog box.

Step 2 In the DHCP Server field, enter either the IP4 address or hostname or the IPv6 address or hostname of the external DHCP server to which DHCP requests are forwarded. Alternatively, click the ellipses to display the Browse DHCP Server dialog box.

Step 3 Double-click a DHCP server from the list, then click **OK** to close the Browse DHCP Server dialog box.

The newly added external DHCP server appears in the list.

Step 4 Choose the interface through which DHCP requests are forwarded to the external DHCP server from the drop-down list.

Step 5 Click **OK** to close the Add Global Relay Servers dialog box.

The newly added global DHCP relay server appears in the list.

Step 6 To modify global DHCP relay server settings, click **Edit** to display the Edit DHCP Global Relay Servers dialog box.

Step 7 Make the desired changes, then click **OK** to close the Edit DHCP Global Relay Servers dialog box. The updated global DHCP relay server settings appear in the list.

Step 8 To change the timeout value, enter the number of seconds for IPv4 or IPv6.

Step 9 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Additional References

For additional information related to implementing DHCPv6, see the following section:

- RFCs, page 23-9

RFCs

RFC	Title
2132	DHCP Options and BOOTP Vendor Extensions
2462	IPv6 Stateless Address Autoconfiguration
5510	DHCP for IPv6

DHCP Monitoring

To monitor DHCP, perform one or more of the following steps:

Path	Purpose
Tools > Command Line Interface Enter the **show running-config dhcpd** command, then click **Send**.	Shows the current DHCP configuration.
Tools > Command Line Interface Enter the **show running-config dhcprelay** command, then click **Send**.	Shows the current DHCP relay service status.
Tools > Command Line Interface Enter the **show ipv6 dhcprelay binding** command, then click **Send**.	Shows the relay binding entries that were created by the relay agent.
Tools > Command Line Interface Enter the **show ipv6 dhcprelay statistics** command, then click **Send**.	Shows DHCP relay agent statistics for IPv6.
Tools > Command Line Interface Enter the **clear config ipv6 dhcprelay** command, then click **Send**.	Clears the IPv6 DHCP relay configuration.
Monitoring > Interfaces > DHCP > DHCP Client Lease Information	Shows configured DHCP client IP addresses.
Monitoring > Interfaces > DHCP > DHCP Server Table	Shows configured dynamic DHCP client IP addreses.
Monitoring > Interfaces > DHCP > DHCP Statistics	Shows DHCP message types, counters, values, directions, messages received, and messages sent.

Feature History for DHCP

Table 23-2 each feature change and the platform release in which it was implemented.

ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 23-2 Feature History for DHCP

Feature Name	Releases	Description
DHCP	7.0(1)	The ASA can provide a DHCP server or DHCP relay services to DHCP clients attached to ASA interfaces. We introduced the following screens: Configuration > Device Management > DHCP > DHCP Relay. Configuration > Device Management > DHCP > DHCP Server.
DHCP for IPv6 (DHCPv6)	9.0(1)	Support for IPv6 was added. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.

PART 4

Configuring Objects and ACLs

Configuring Objects

This chapter describes how to configure reusable named objects and groups for use in your configuration, and it includes the following sections:

- Information About Objects, page 25-1
- Licensing Requirements for Objects, page 25-1
- Configuring Objects, page 25-2
- Monitoring Objects, page 25-16
- Feature History for Objects, page 25-16

Information About Objects

Objects are reusable components for use in your configuration. They can be defined and used in ASA configurations in the place of inline IP addresses, services, names, and so on. Objects make it easy to maintain your configurations because you can modify an object in one place and have it be reflected in all other places that are referencing it. Without objects you would have to modify the parameters for every feature when required, instead of just once. For example, if a network object defines an IP address and subnet mask, and you want to change the address, you only need to change it in the object definition, not in every feature that refers to that IP address.

Licensing Requirements for Objects

Model	License Requirement
All models	Base License.

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

- Supports IPv6.
- The ASA does not support IPv6 nested network object groups, so you cannot group an object with IPv6 entries under another IPv6 object group.
- You can mix IPv4 and IPv6 entries in a network object group; you cannot use a mixed object group for NAT.

Additional Guidelines and Limitations

- Object must have unique names. While you might want to create a network object group named "Engineering" and a service object group named "Engineering," you need to add an identifier (or "tag") to the end of at least one object group name to make it unique. For example, you can use the names "Engineering_admins" and "Engineering_hosts" to make the object group names unique and to aid in identification.
- Objects and object groups share the same name space.
- You cannot remove an object or make an object empty if it is used in a command.

Configuring Objects

- Configuring Network Objects and Groups, page 25-2
- Configuring Service Objects and Service Groups, page 25-4
- Configuring Local User Groups, page 25-7
- Configuring Security Group Object Groups, page 25-8
- Configuring Regular Expressions, page 25-10
- Configuring Time Ranges, page 25-15

Note For other objects not included in this chapter, see the following chapters:

- Local Users—See Chapter 46, "Configuring AAA Servers and the Local Database."
- Class Maps—See Chapter 57, "Getting Started with Application Layer Protocol Inspection."
- Inspect Maps—See Chapter 57, "Getting Started with Application Layer Protocol Inspection."
- TCP Maps—See the "Configuring Connection Settings" section on page 71-5.

Configuring Network Objects and Groups

This section describes how to configure network objects and groups, and it includes the following topics:

- Configuring a Network Object, page 25-3
- Configuring a Network Object Group, page 25-3

Configuring a Network Object

A network object can contain a host, a network IP address, or a range of IP addresses, a fully qualified domain name (FQDN). You can also enable NAT rules on the object (excepting FQDN objects). (See Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)," for more information.)

Detailed 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—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must contain 64 characters or fewer.
- Type—Either Network, Host, Range, or FQDN.
- 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. If you select Range as the object type, the IP Address field changes to allow you to enter a Start Address and an End address.
- 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. (This field is not available if you enter the object type as Host.)
- Description—(Optional) The description of the network object (up to 200 characters in length).

Note To add NAT rules to the network object, see Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)," for more information.

Step 4 Click **OK**.

Step 5 Click **Apply** to save the configuration.

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

Configuring a Network Object Group

Network object groups can contain multiple network objects as well as inline networks. Network object groups can support a mix of both IPv4 and IPv6 addresses.

Restrictions

You cannot use a mixed IPv4 and IPv6 object group for NAT, or object groups that include FQDN objects.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Network Objects/Groups**.

Step 2 Click click **Add > Network Object Group** to add either a new object or a new object group.

You can also add or edit network object groups from the Addresses side pane in a rules window, or when you add 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 Network Object Group dialog box appears.

Step 3 In the Group Name field, enter a group name.

Use characters a to z, A to Z, 0 to 9, a period, a comma, a dash, or an underscore. The name must contain 64 characters or fewer.

Step 4 (Optional) In the Description field, enter a description, up to 200 characters in length.

Step 5 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 the object in the Members in Group pane, or select the object and click **Remove**.

Step 6 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.

Configuring Service Objects and Service Groups

Service objects and groups identify protocols and ports. This section describes how to configure service objects, service groups, TCP and UDP port service groups, protocol groups, and ICMP groups, and it includes the following topics:

- Configuring a Service Object, page 25-5
- Configuring a Service Group, page 25-5
- Configuring a TCP or UDP Port Service Group, page 25-6

- Configuring an ICMP Group, page 25-6
- Configuring an ICMP Group, page 25-6

Configuring a Service Object

The service object can contain a protocol, ICMP, ICMPv6, TCP or UDP port or port ranges.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Service Object/Group**.

Step 2 Choose **Add > Service Object** from the drop-down list.

Step 3 In the name field, enter a name for the service object. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or fewer.

Step 4 From the Service Type field, choose the desired type: tcp, udp, icmp, or icmp6 protocol.

Step 5 (Optional) If you chose tcp or udp as the Service Type, enter the following:

- Destination Port/Range
- Source Port/Range—Lists the protocol source ports/ranges.
- Description—Lists the service group description.

Step 6 (Optional) If you chose icmp or icmp6 as the Service Type, enter the following:

- ICMP Type—Lists the service group ICMP type.
- ICMP Code—(Optional) Valid values range from 1 to 255.
- Description—(Optional) Lists the service group description.

Step 7 If you chose protocol as the Service Type, enter the following:

- Protocol—Lists the service group protocol.
- Description—(Optional) Lists the service group description.

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

Configuring a Service Group

A service object group includes a mix of protocols, if desired, including optional source and destination ports for TCP or UDP.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Service Object/Group**.

Step 2 Choose **Add > Service Group** from the drop-down list.

The Add Service Group dialog box appears.

Step 3 In the Name field, enter a name for the new service group. The name can be up to 64 characters in length and must be unique for all object groups. A service group name cannot share a name with other objects and groups.

Step 4 In the Description field, enter a description for this service group (up to 200 characters in length).

Step 5 To add an existing service object or group, or predefined protocol or port, click the **Existing Service/Service Group** radio button, select the entry from the Name field, and click **Add**.

Step 6 To create a new service, click the **Create new member** radio button and then choose the Service Type from the drop-down list:

- If you choose tcp, udp, or tcp/udp, enter a name, the destination port/range, the source port/range, and an optional description.
- If you choose icmp or icmp6, enter a name, the ICMP Type (from the Existing Service/Service Group list), an ICMP Code (a value from 0-255), and an optional description.
- If you choose protocol, enter a name, the protocol, and an optional description.

Click **Add** to add the new service.

Step 7 Click **OK**, and then **Apply**.

Configuring a TCP or UDP Port Service Group

A TCP or UDP service group includes a group of ports for a specific protocol (TCP, UDP, or TCP-UDP).

Step 1 Choose **Configuration > Firewall > Objects > Service Object/Group**.

Step 2 Choose **Add > TCP Service Group**, **UDP Service Group**, or **TCP-UDP Service Group** from the drop-down list.

The Add Service Group dialog box appears.

Step 3 In the Name field, enter a name for the new service group. The name can be up to 64 characters in length and must be unique for all object groups. A service group name cannot share a name with other objects and groups.

Step 4 In the Description field, enter a description for this service group (up to 200 characters in length).

Step 5 To add an existing service group, or predefined port, click the **Existing Service/Service Group** radio button, select the entry from the Name field, and click **Add**.

Step 6 To create a new port, click the **Create new member** radio button, enter the port name, number, or range and then click **Add** to add the new port.

Step 7 Click **OK**, and then **Apply**.

Configuring an ICMP Group

An ICMP group includes multiple ICMP types.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Service Object/Group**.

Step 2 Choose **Add > ICMP Group** from the drop-down list.

The Add ICMP Group dialog box appears.

Step 3 In the Name field, enter a name for the new ICMP group. The name can be up to 64 characters in length and must be unique for all object groups. An ICMP group name cannot share a name with other objects and groups..

Step 4 In the Description field, enter a description for this ICMP group (up to 200 characters in length).

Step 5 To add an existing ICMP group, or predefined type, click the **Existing Service/Service Group** radio button, select the entry from the Name field, and click **Add**.

Step 6 To create a new type, click the **Create new member** radio button, enter the type name or number, and then click **Add** to add the new type.

Step 7 Click **OK**, and then **Apply**.

Configuring a Protocol Group

A protocol group contains IP protocol types.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Service Object/Group**.

Step 2 Choose **Add > Protocol Group** from the drop-down list.

The Add Protocol Group dialog box appears.

Step 3 In the Name field, enter a name for the new group. The name can be up to 64 characters in length and must be unique for all object groups. A group name cannot share a name with other objects and groups.

Step 4 In the Description field, enter a description for this group (up to 200 characters in length).

Step 5 To add an existing protocol group, or predefined protocol, click the **Existing Service/Service Group** radio button, select the entry from the Name field, and click **Add**.

Step 6 To create a new protocol, click the **Create new member** radio button, enter the protocol name or number, and then click **Add** to add the new protocol.

Step 7 Click **OK**, and then **Apply**.

Configuring Local User Groups

You can create local user groups for use in features that support the identity firewall (IDFW) by including the group in an extended ACL, which in turn can be used in an access rule, for example.

The ASA sends an LDAP query to the Active Directory server for user groups globally defined in the Active Directory domain controller. The ASA imports these groups for identity-based rules. However, the ASA might have localized network resources that are not defined globally that require local user groups with localized security policies. Local user groups can contain nested groups and user groups that are imported from Active Directory. The ASA consolidates local and Active Directory groups.

A user can belong to local user groups and user groups imported from Active Directory.

Prerequisites

See Chapter 47, “Configuring the Identity Firewall,” to enable IDFW.

Detailed Steps

Step 1 Open the **Configuration > Firewall > Objects > Local User Groups** pane.

A table of user groups and their members appears.

Step 2 To add a group, click **Add**. The Add User Object Group dialog appears.

Step 3 Enter a name and description for the group.

The group name can contain any character including [a-z], [A-Z], [0-9], [!@#$%^&()-_{}.]. If the group name contains a space, you must enclose the name in quotation marks.

Step 4 From the Domain list, select the default domain for users in this group or click **Manage** to add a new domain or edit and existing domain.

Step 5 To add existing groups to this group, enter a search string in the text box and click **Find**.

Step 6 To add users to the group, enter a search string in the text box and click **Find**.

Step 7 Select groups and click the **Add** button to add them to the group.

Step 8 Select users and click the **Add** button to add them to the group.

Step 9 Click **OK** to save your changes.

Configuring Security Group Object Groups

You can create security group object groups for use in features that support Cisco TrustSec by including the group in an extended ACL, which in turn can be used in an access rule, for example.

When integrated with Cisco TrustSec, the ASA downloads security group information from the ISE. The ISE acts as an identity repository, by providing Cisco TrustSec tag to user identity mapping and Cisco TrustSec tag to server resource mapping. You provision and manage security group access lists centrally on the ISE.

However, the ASA might have localized network resources that are not defined globally that require local security groups with localized security policies. Local security groups can contain nested security groups that are downloaded from the ISE. The ASA consolidates local and central security groups.

To create local security groups on the ASA, you create a local security object group. A local security object group can contain one or more nested security object groups or Security IDs or security group names. User can also create a new Security ID or security group name that does not exist on the ASA.

You can use the security object groups you create on the ASA to control access to network resources. You can use the security object group as part of an access group or service policy.

Prerequisites

See Chapter 48, “Configuring the ASA to Integrate with Cisco TrustSec,” to enable TrustSec.

Detailed Steps

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Objects > Security Group Object Groups**. The Security Group Object Groups pane appears:

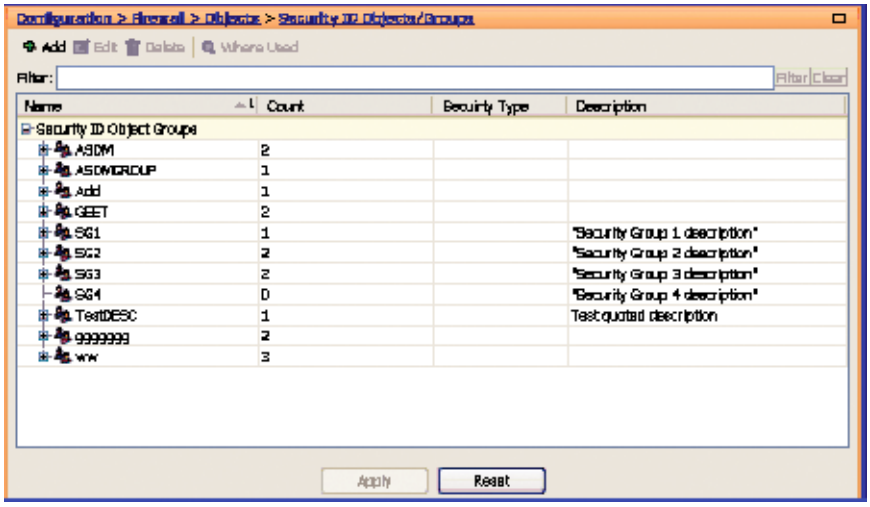

The Security Group Object Groups pane lists the members of the security object group and shows the number of members in the Count column. Click **Where Used** to display where the selected security group object is used in an access list or nested in another security group object.

Step 2 Click **Add**. The Add Security ID Object Group dialog box appears.

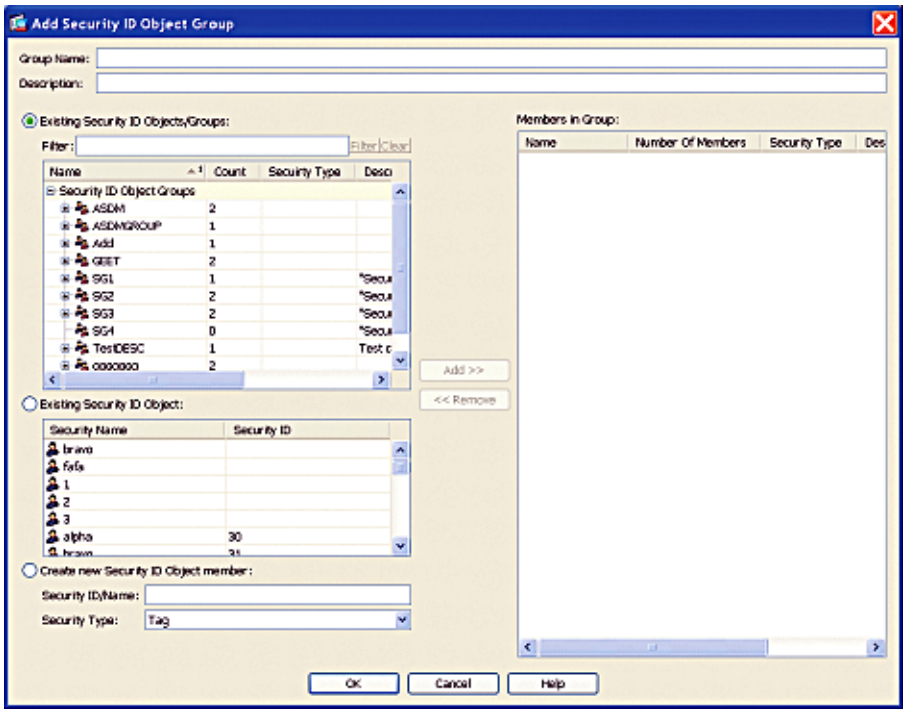

Step 3 In the Group Name field, enter the name for the group as a 32-byte case sensitive string. The group name can contain any character including [a-z], [A-Z], [0-9], [!@#$%^&()-_{}.].

Step 4 In the Description field, enter a description for the group.

Step 5 Add members to the security group object by performing the following task:

a. Select one of the following options:

- Existing Security ID Objects/Groups radio button
- Existing Security ID Object radio button

  In the Filter field, enter the security object ID number or the name of the security group and click **Filter**. Use wildcards to broaden the search for security groups.

b. Click **Add** to select it as Members in Group.

A security object group must contain at least one member.

c. Continue selecting members and clicking Add. You can create nested security object groups by selecting existing security ID object/groups and existing security ID objects.

Step 6 Create a locally defined object by performing the following tasks:

a. Click the **Create new Security ID Object member** radio button.

b. From the Security Type drop-down field, select Tag or Name.

An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB) by the ISE. Security group names are created on the ISE and provide user-friendly names for security groups. The security group table maps SGTs to security group names.

c. In the Security ID/Name field, enter a number from 1 to 65533 for A Tag security type or a 32-byte case-sensitive string for a Name security type.

A security group has a single name assigned to it. The same name can only be associated with a single SGT.

Step 7 Click OK. The Security ID Objects/Groups pane reappears.

Step 8 Click **Apply** to save the changes to the running configuration.

Configuring Regular Expressions

- Creating a Regular Expression, page 25-10
- Creating a Regular Expression Class Map, page 25-14

Creating a Regular Expression

A regular expression matches text strings either literally as an exact string, or by using *metacharacters* so that 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 a URL string inside an HTTP packet.

Guidelines

Note As an optimization, the ASA 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 25-1 lists the metacharacters that have special meanings.

Table 25-1 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, and so on.
+	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.

Table 25-1 regex Metacharacters (continued)

Character	Description	Notes
\	Escape character	When used with a metacharacter, matches a literal character. For example, \[matches the left square bracket.
char	Character	When character is not a metacharacter, matches the literal character.
\r	Carriage return	Matches a carriage return 0x0d.
\n	Newline	Matches a new line 0x0a.
\t	Tab	Matches a tab 0x09.
\f	Formfeed	Matches a form feed 0x0c.
\x*NN*	Escaped hexadecimal number	Matches an ASCII character using hexadecimal (exactly two digits).
NNN	Escaped octal number	Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space.

Detailed Steps

Step 1 Choose **Configuration > Global Objects > Regular Expressions**.

Step 2 In the Regular Expressions area, click **Add**.

The Add Regular Expression dialog box appears.

Step 3 In the Name field, name the expression, up to 40 characters in length.

Step 4 (Optional) Click **Build** to use the Creating a Regular Expression Class Map dialog box. See Table 25-1 on page 25-11 for more information about metacharacters.

- 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.

Step 5 If you do not use the Build tool, enter the regular expression manually in the Value field, up to 100 characters in length. Refer to the metacharacters in Table 25-1.

Step 6 To test the regular expression before adding it, click **Test**.

The Test Regular Expression dialog box appears.

- 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.

Creating a Regular Expression Class Map

A regular expression class map identifies one or more regular expressions. You can use a regular expression class map to match the content of certain traffic; for example, you can match URL strings inside HTTP packets.

Prerequisites

Create one or more regular expressions according to the "Creating a Regular Expression" section on page 25-10.

Detailed Steps

Step 1 Choose **Configuration > Global Objects > Regular Expressions**.

Step 2 In the Regular Expression Classes area, click **Add**.

- 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.

Configuring Time Ranges

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 ASA.

A time range consists of a start time, an end time, and optional recurring entries.

Guidelines

- Multiple periodic entries are allowed per time range. If a time range has both absolute and periodic values specified, then the periodic values are evaluated only after the absolute start time is reached, and they are not further evaluated after the absolute end time is reached.
- Creating a time range does not restrict access to the device. This procedure defines the time range only.

Detailed Steps

Step 1 Choose **Configuration > Global Objects > Time Ranges**.

Step 2 Click **Add**.

The Add Time Range window appears.

Step 3 In the Time Range Name field, enter a time range name, with no spaces.

Step 4 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 5 (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, in the Recurring Time Ranges area, click **Add.**

The Add Recurring Time Range dialog box appears.

Step 6 Do one of the following:

- 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**.
- 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 7 Click **OK**, and then click **Apply**.

Monitoring Objects

To view which 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.

Feature History for Objects

Table 25-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 25-2 Feature History for Object Groups

Feature Name	Platform Releases	Feature Information
Object groups	7.0(1)	Object groups simplify access list creation and maintenance.
Objects	8.3(1)	Object support was introduced.
User Object Groups for Identity Firewall	8.4(2)	User object groups for identity firewall were introduced.
Mixed IPv4 and IPv6 network object groups	9.0(1)	Previously, network object groups could only contain all IPv4 addresses or all IPv6 addresses. Now network object groups can support a mix of both IPv4 and IPv6 addresses. **Note** You cannot use a mixed object group for NAT.
Security Group Object Groups for Cisco TrustSec	8.4(2)	Security group object groups for TrustSec were introduced.
Extended ACLand object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule

Using the ACL Manager

This chapter describes how to configure extended access lists (also known as access control lists), and it includes the following sections:

- Information About the ACL Manager, page 26-1
- Licensing Requirements for the ACL Manager, page 26-1
- Guidelines and Limitations, page 26-2
- Adding ACLs and ACEs, page 26-2
- Feature History for the ACL Manager, page 26-5

Information About the ACL Manager

Access control lists (ACLs) are used to control network access or to specify traffic for many features to act upon. An ACL is made up of one or more access control entries (ACEs) in which you can specify the line number to insert the ACE, the source and destination addresses, and, depending upon the ACE type, the protocol, the ports (for TCP or UDP), or the ICMP type.

The ACL Manager dialog box lets you define 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 ASA 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 ASA denies it. ACEs are referred to as rules in this section.

For information about adding ACLs and ACEs, see the "Adding ACLs and ACEs" section on page 26-2.

For information about finding specific ACLs and ACEs in your configuration, see the "Using the Find Function in the ACL Manager Pane" section on page 4-15.

Licensing Requirements for the ACL Manager

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 modes only.

IPv6 Guidelines

IPv6 is supported.

Additional Guidelines and Limitations

The following guidelines and limitations apply to creating an extended access list:

- Enter the access list name in uppercase letters so that the name is easy to see in the configuration. You might want to name the access list for the interface (for example, INSIDE), or you can name it for the purpose for which it is created (for example, NO_NAT or VPN).
- You can specify the source and destination ports only for the TCP or UDP protocols. For a list of permitted keywords and well-known port assignments, see the "TCP and UDP Ports" section on page A-11. DNS, Discard, Echo, Ident, NTP, RPC, SUNRPC, and Talk each require one definition for TCP and one for UDP. TACACS+ requires one definition for port 49 on TCP.

Adding ACLs and ACEs

An access list (ACL) is made up of one or more access list entries (ACEs). To create an ACL, you start by creating an ACE and applying a list name. An ACL with one entry is still considered a list, although you can add multiple ACEs to the list.

To add an ACL and then add an ACE to that ACL, perform the following steps:

Step 1 Choose **Configuration > Firewall > Advanced > ACL Manager**.

Step 2 Select **Add > Add ACL**. Adds an ACL configurable for IPv4 or IPv6 traffic.

Step 3 In the ACL name field, add a descriptive name for the ACL, and click **OK**.

Your newly created ACL appears in the window.

Step 4 Select the newly created ACL, click **Add**, and from the drop-down list, choose **Add ACE**.

Step 5 In the Action field of the Add ACE window, 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 6 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.

IPv6 must be enabled on at least one interface before you can configure an ACE with an IPv6 address. For more information about enabling IPv6 on an interface, see the "Configuring IPv6 Addressing" section on page 17-15.

Step 7 Select a destination to specify the IP addresses (host or network) that are permitted or denied to send traffic to the IP addresses listed in the Source section.

Step 8 Specify the service to which this ACE applies. You can type a known service into the window or click browse to select from a list of services.

Service groups let you identify multiple non-contiguous port numbers that you want 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—Enables you to 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—Enables you to 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.

Step 9 (Optional) Add text that provides a brief description of this rule. A description line can be up to 100 characters long, yet you can break a description into multiple lines.

Note If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Step 10 (Optional) Check the Enable Logging check box to enable or disable logging or specify the use of the default logging settings. If logging is enabled, the Syslog Level and Log Interval fields become available.

a. If logging is enables, choose a logging level to specify logging activity. The default is Informational. For information about logging levels, see the "Severity Levels" section on page 96-3.

b. Choose a logging interval to display the interval, in seconds, that is used to limit how many messages at this logging level can be sent.

Step 11 Set the source service (TCP, UDP, and TCP/UDP only).

Step 12 Set the logging interval to establish the number of seconds between log messages. The default is 300.

Step 13 Set the time range during which the rule is applied.

Step 14 Click **Apply** to save the ACL and ACE to the running configuration.

To see a condensed view of all ACLs in your configuration, click **Collapse All** below the ACL Manager window. To see a comprehensive view of all ACLs and ACEs in your configuration, click **Expand All**.

For information about finding specific ACLs and ACEs in your configuration, see the "Using the Find Function in the ACL Manager Pane" section on page 4-15.

Using Standard ACLs in the ACL Manager

Standard ACLs identify the destination IP addresses (not source addresses). Standard ACLs cannot be applied to interfaces to control traffic.

To add a standard ACL to your configuration, perform the following steps:

Step 1 Click **Add**, and from the drop-down list, choose **Add ACL**.

Step 2 In the Add ACL dialog box, add a name or number (without spaces) to identify the ACL.

Step 3 Click **OK**

The ACL name appears in the main pane.

Step 4 Select the newly created ACL, click **Add**, and from the drop-down list, choose **Add ACE**.

The Add ACE dialog box appears.

Step 5 (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 6 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 7 In the Address field, enter the IP address of the destination to which you want to perform 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 8 (Optional) In the Description field, enter a description that makes an ACE easier to understand.

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

Note If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Step 9 Click **OK**.

The newly created ACE appears under the ACL.

Step 10 Click **Apply** to save the ACE to your configuration.

Feature History for the ACL Manager

Table 26-1 lists the release history for this feature.

Table 26-1 Feature History for Extended Access Lists

Feature Name	Releases	Feature Information
Extended access lists	7.0(1)	Access lists are used to control network access or to specify traffic for many features to act upon. An extended access control list is made up of one or more access control entries (ACEs) in which you can specify the line number to insert the ACE, the source and destination addresses, and, depending upon the ACE type, the protocol, the ports (for TCP or UDP), or the IPCMP type (for ICMP).

Adding a Standard Access Control List

This chapter describes how to configure a standard ACL and includes the following sections:

- Information About Standard ACLs, page 27-1
- Licensing Requirements for Standard ACLs, page 27-1
- Guidelines and Limitations, page 27-1
- Default Settings, page 27-2
- Adding Standard ACLs, page 27-3
- Feature History for Standard ACLs, page 27-4

Information About Standard ACLs

Standard access lists identify the destination IP addresses of OSPF routes and can be used in a route map for OSPF redistribution. Standard access lists cannot be applied to interfaces to control traffic.

Licensing Requirements for Standard ACLs

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, page 27-2
- Firewall Mode Guidelines, page 27-2
- IPv6 Guidelines, page 27-2
- Additional Guidelines and Limitations, page 27-2

Context Mode Guidelines

Supported in single context mode only.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

The following guidelines and limitations apply for 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.
- When specifying a source, local, or destination address, use the following guidelines:
  - Use a 32-bit quantity in four-part, dotted-decimal format.
- If you add descriptive remarks to your ACL with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Default Settings

Table 27-1 lists the default settings for standard ACL parameters.

Table 27-1 Default Standard Access List Parameters

Parameters	Default
deny	The ASA denies all packets on the originating interface unless you specifically permit access. Access list logging generates system log message 106023 for denied packets. Deny packets must be present to log denied packets.

Adding Standard ACLs

This section includes the following topics:

- Using Standard ACLs, page 27-3

Using 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.

This section includes the following topics:

- Adding a Standard ACL, page 27-3
- Adding an ACE to a Standard ACL, page 27-3
- Editing an ACE in a Standard ACL, page 27-4

Adding a Standard ACL

To add a standard ACL to your configuration, 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 27-3.

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 27-3. For information about editing ACEs, see the "Editing an ACE in a Standard ACL" section on page 27-4

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 perform 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 an 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 existing 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**.

Feature History for Standard ACLs

Table 27-2 lists the release history for this feature.

Table 27-2 Feature History for Standard Access Lists

Feature Name	Releases	Feature Information
Standard ACLs	7.0(1)	Standard ACLs identify the destination IP addresses of OSPF routes, which can be used in a route map for OSPF redistribution. The feature was introduced.

Adding a Webtype Access Control List

Webtype ACLs are added to a configuration that supports filtering for clientless SSL VPN. This chapter describes how to add an ACLto the configuration that supports filtering for WebVPN.

This chapter includes the following sections:

- Licensing Requirements for Webtype ACLs, page 28-1
- Guidelines and Limitations, page 28-1
- Default Settings, page 28-2
- Using Webtype ACLs, page 28-2
- Feature History for Webtype Access Lists, page 28-6
- Feature History for Webtype Access Lists, page 28-6

Licensing Requirements for Webtype ACLs

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, page 28-1
- Firewall Mode Guidelines, page 28-1
- Additional Guidelines and Limitations, page 28-2

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

The following guidelines and limitations apply to Webtype ACLs:

- Smart tunnel ACEs filter on a per-server basis only, so you cannot create smart tunnel ACEs to permit or deny access to directories or to permit or deny access to specific smart tunnel-enabled applications.
- If you add descriptive remarks about your ACL with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Default Settings

Table 28-1 lists the default settings for Webtype access lists parameters.

Table 28-1 Default Webtype Access List Parameters

Parameters	Default
deny	The ASA denies all packets on the originating interface unless you specifically permit access.
log	Access list logging generates system log message 106023 for denied packets. Deny packets must be present to log denied packets.

Using Webtype ACLs

This section includes the following topics:

- Adding a Webtype ACL and ACE, page 28-3
- Editing Webtype ACLs and ACEs, page 28-4
- Deleting Webtype ACLs and ACEs, page 28-5

Task Flow for Configuring Webtype ACLs

Use the following guidelines to create and implement an ACL:

- Create an ACL by adding an ACE and applying an ACL name. See the "Using Webtype ACLs" section on page 28-2.
- Apply the ACL to an interface. See the "Configuring Access Rules" section on page 51-7 for more information.

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 on a per-server basis only, so you cannot create smart tunnel ACEs to permit or deny access to directories or to permit or deny access to specific smart tunnel-enabled applications.

To configure a webtype ACL, 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>

Wildcard 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.

Wildcard characters can be used in the with regular expression in the address 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 permit a range of IP addresses from 10.2.2.20 through 10.2.2.31, enter the following: **10.2.2.[20-31]**

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.

b. The Browse Time Range dialog box appears.

c. Click **Add**.

d. The Add Time Range dialog box appears.

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

f. Enter the Start Time and the End Time.

g. 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 ACLs, you can click the following radio buttons to filter which ACLs appear in the main pane: IPv4 andIPv6, IPv4 only, or IPv6 Only.

Editing Webtype ACLs and ACEs

To edit a webtype ACL or ACT, perform the following steps:

Step 1 Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Web ACLs**.

Step 2 Choose the ACL type to edit by clicking one of the following radio buttons:

- **IPv4 and IPv6**— Shows ACLs that have both IPv4 and IPv6 addresses only.
- **IPv4 Only**—Shows ACLs 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 28-3.

Step 4 Click **OK**.

Step 5 Click **Apply** to save the changes to your configuration.

Deleting Webtype ACLs and 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 ACL type to edit by clicking one of the following radio buttons:

- **IPv4 and IPv6**— Shows ACLs that have both IPv4 and IPv6 addresses only.
- **IPv4 Only**—Shows ACLs 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 delete.

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**.

The selected items are removed from the viewing pane.

Note If you deleted an item in error and want to restore it to your configuration, click **Reset** before you click Apply. The deleted item reappears in the viewing pane.

Step 5 Click **Apply** to save the change to the configuration.

Feature History for Webtype Access Lists

Table 28-2 lists the release history for this feature.

Table 28-2 Feature History for Webtype Access Lists

Feature Name	Releases	Feature Information
Webtype access lists	7.0(1)	Webtype ACLs are access lists that are added to a configuration that supports filtering for clientless SSL VPN. We introduced the feature.
Unified ACL for IPv4 and IPv6	9.0(1)	ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options

PART 5

Configuring IP Routing

Routing Overview

This chapter describes underlying concepts of how routing behaves within the ASA, and the routing protocols that are supported.

This chapter includes the following sections:

- Information About Routing, page 30-1
- How Routing Behaves Within the ASA, page 30-4
- Supported Internet Protocols for Routing, page 30-5
- Information About the Routing Table, page 30-5
- Disabling Proxy ARPs, page 30-11

Information About Routing

Routing is the act of moving information across an internetwork from a source to a destination. Along the way, at least one intermediate node typically is encountered. Routing involves two basic activities: determining optimal routing paths and transporting information groups (typically called packets) through an internetwork. In the context of the routing process, the latter of these is referred to as packet switching. Although packet switching is relatively straightforward, path determination can be very complex.

This section includes the following topics:

- Switching, page 30-1
- Path Determination, page 30-2
- Supported Route Types, page 30-2

Switching

Switching algorithms is relatively simple; it is the same for most routing protocols. In most cases, a host determines that it must send a packet to another host. Having acquired a router address by some means, the source host sends a packet addressed specifically to a router physical (Media Access Control [MAC]-layer) address, this time with the protocol (network layer) address of the destination host.

As it examines the packet destination protocol address, the router determines that it either knows or does not know how to forward the packet to the next hop. If the router does not know how to forward the packet, it typically drops the packet. If the router knows how to forward the packet, however, it changes the destination physical address to that of the next hop and transmits the packet.

The next hop may be the ultimate destination host. If not, the next hop is usually another router, which executes the same switching decision process. As the packet moves through the internetwork, its physical address changes, but its protocol address remains constant.

Path Determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which include route information. Route information varies depending on the routing algorithm used.

Routing algorithms fill routing tables with a variety of information. Destination or next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with a next hop.

Routing tables also can include other information, such as data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

Note Asymmetric routing is only supported for Active/Active failover in multiple context mode. For more information, see the "Configuring Active/Active Failover" section on page 13-8.

Supported Route Types

There are several route types that a router can use. The ASA uses the following route types:

- Static Versus Dynamic, page 30-3
- Single-Path Versus Multipath, page 30-3
- Flat Versus Hierarchical, page 30-3
- Link-State Versus Distance Vector, page 30-3

Static Versus Dynamic

Static routing algorithms are hardly algorithms at all, but are table mappings established by the network administrator before the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a router to which all unroutable packets are sent), for example, can be designated to act as a repository for all unroutable packets, ensuring that all messages are at least handled in some way.

Single-Path Versus Multipath

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multipath algorithms permit traffic multiplexing over multiple lines. The advantages of multipath algorithms are substantially better throughput and reliability, which is generally called load sharing.

Flat Versus Hierarchical

Some routing algorithms operate in a flat space, while others use routing hierarchies. In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from nonbackbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more nonbackbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas. In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intradomain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

Link-State Versus Distance Vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state

algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors. Typically, link-state algorithms are used in conjunction with OSPF routing protocols.

How Routing Behaves Within the ASA

The ASA uses both routing table and XLATE tables for routing decisions. To handle destination IP translated traffic, that is, untranslated traffic, the ASA searches for existing XLATE, or static translation to select the egress interface.

This section includes the following topics:

- Egress Interface Selection Process, page 30-4
- Next Hop Selection Process, page 30-4

Egress Interface Selection Process

The selection process follows these steps:

1. If a destination IP translating XLATE already exists, the egress interface for the packet is determined from the XLATE table, but not from the routing table.
2. If a destination IP translating XLATE does not exist, but a matching static translation exists, then the egress interface is determined from the static route and an XLATE is created, and the routing table is not used.
3. If a destination IP translating XLATE does not exist and no matching static translation exists, the packet is not destination IP translated. The ASA processes this packet by looking up the route to select the egress interface, then source IP translation is performed (if necessary).

   For regular dynamic outbound NAT, initial outgoing packets are routed using the route table and then creating the XLATE. Incoming return packets are forwarded using existing XLATE only. For static NAT, destination translated incoming packets are always forwarded using existing XLATE or static translation rules.

Next Hop Selection Process

After selecting the egress interface using any method described previously, an additional route lookup is performed to find out suitable next hop(s) that belong to a previously selected egress interface. If there are no routes in the routing table that explicitly belong to a selected interface, the packet is dropped with a level 6 syslog message 110001 generated (no route to host), even if there is another route for a given destination network that belongs to a different egress interface. If the route that belongs to a selected egress interface is found, the packet is forwarded to the corresponding next hop.

Load sharing on the ASA is possible only for multiple next hops available using a single egress interface. Load sharing cannot share multiple egress interfaces.

If dynamic routing is in use on the ASA and the route table changes after XLATE creation (for example, route flap), then destination translated traffic is still forwarded using the old XLATE, not via the route table, until XLATE times out. It may be either forwarded to the wrong interface or dropped with a level 6 syslog message 110001 generated (no route to host), if the old route was removed from the old interface and attached to another one by the routing process.

The same problem may happen when there are no route flaps on the ASA itself, but some routing process is flapping around it, sending source-translated packets that belong to the same flow through the ASA using different interfaces. Destination-translated return packets may be forwarded back using the wrong egress interface.

This issue has a high probability in some security traffic configurations, where virtually any traffic may be either source-translated or destination-translated, depending on the direction of the initial packet in the flow. When this issue occurs after a route flap, it can be resolved manually by using the **clear xlate** command, or automatically resolved by an XLATE timeout. The XLATE timeout may be decreased if necessary. To ensure that this issue rarely occurs, make sure that there are no route flaps on the ASA and around it. That is, ensure that destination-translated packets that belong to the same flow are always forwarded the same way through the ASA.

Supported Internet Protocols for Routing

The ASA supports several Internet protocols for routing. Each protocol is briefly described in this section.

- Enhanced Interior Gateway Routing Protocol (EIGRP)

  EIGRP is a Cisco proprietary protocol that provides compatibility and seamless interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into Enhanced IGRP, and vice versa, so it is possible to add Enhanced IGRP gradually into an existing IGRP network.

  For more information about configuring EIGRP, see the "Configuring EIGRP" section on page 36-4.

- Open Shortest Path First (OSPF)

  OSPF is a routing protocol developed for Internet Protocol (IP) networks by the interior gateway protocol (IGP) working group of the Internet Engineering Task Force (IETF). OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area includes an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

  For more information about configuring OSPF, see the "Configuring OSPFv2" section on page 33-5.

- Routing Information Protocol (RIP)

  RIP is a distance-vector protocol that uses hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system.

  For more information about configuring RIP, see the "Configuring RIP" section on page 34-4.

Information About the Routing Table

This section includes the following topics:

- Displaying the Routing Table, page 30-6
- How the Routing Table Is Populated, page 30-6
- How Forwarding Decisions Are Made, page 30-8
- Dynamic Routing and Failover, page 30-8
- Dynamic Routing and Clustering, page 30-9

- Dynamic Routing in Multiple Context Mode, page 30-10

Displaying the Routing Table

To show all routes in ASDM that are in the routing table, choose **Monitoring > Routing > Routes**.

In this pane, each row represents one route.

How the Routing Table Is Populated

The ASA routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the RIP, EIGRP, and OSPF routing protocols. Because the ASA can run multiple routing protocols in addition to having static and connected routes in the routing table, it is possible that the same route is discovered or entered in more than one manner. When two routes to the same destination are put into the routing table, the one that remains in the routing table is determined as follows:

- If the two routes have different network prefix lengths (network masks), then both routes are considered unique and are entered into the routing table. The packet forwarding logic then determines which of the two to use.

  For example, if the RIP and OSPF processes discovered the following routes:

  - RIP: 192.168.32.0/24
  - OSPF: 192.168.32.0/19

  Even though OSPF routes have the better administrative distance, both routes are installed in the routing table because each of these routes has a different prefix length (subnet mask). They are considered different destinations and the packet forwarding logic determines which route to use.

- If the ASA learns about multiple paths to the same destination from a single routing protocol, such as RIP, the route with the better metric (as determined by the routing protocol) is entered into the routing table.

  Metrics are values associated with specific routes, ranking them from most preferred to least preferred. The parameters used to determine the metrics differ for different routing protocols. The path with the lowest metric is selected as the optimal path and installed in the routing table. If there are multiple paths to the same destination with equal metrics, load balancing is done on these equal cost paths.

- If the ASA learns about a destination from more than one routing protocol, the administrative distances of the routes are compared, and the routes with lower administrative distance are entered into the routing table.

Administrative Distances for Routes

You can change the administrative distances for routes discovered by or redistributed into a routing protocol. If two routes from two different routing protocols have the same administrative distance, then the route with the lower *default* administrative distance is entered into the routing table. In the case of EIGRP and OSPF routes, if the EIGRP route and the OSPF route have the same administrative distance, then the EIGRP route is chosen by default.

Administrative distance is a route parameter that the ASA uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. Because the routing protocols have metrics based on algorithms that are different from the other protocols, it is not always possible to determine the best path for two routes to the same destination that were generated by different routing protocols.

Each routing protocol is prioritized using an administrative distance value. Table 30-1 shows the default administrative distance values for the routing protocols supported by the ASA.

Table 30-1 Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
Static route	1
EIGRP Summary Route	5
Internal EIGRP	90
OSPF	110
RIP	120
EIGRP external route	170
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the ASA receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the ASA chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the ASA would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you use the **distance-ospf** command to change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the ASA on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The EIGRP, OSPF, and RIP routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the ASA routing table.

Backup Routes

A backup route is registered when the initial attempt to install the route in the routing table fails because another route was installed instead. If the route that was installed in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route and requests them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes for the failed route, the preferred route is chosen based on administrative distance.

Because of this process, you can create floating static routes that are installed in the routing table when the route discovered by a dynamic routing protocol fails. A floating static route is simply a static route configured with a greater administrative distance than the dynamic routing protocols running on the ASA. When the corresponding route discovered by a dynamic routing process fails, the static route is installed in the routing table.

How Forwarding Decisions Are Made

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.
- If the destination matches more than one entry in the routing table, and the entries all have the same network prefix length, the two entries with identical network prefixes and different interfaces cannot coexist in the routing table.
- If the destination matches more than one entry in the routing table, and the entries have different network prefix lengths, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

For example, a packet destined for 192.168.32.1 arrives on an interface of an ASA with the following routes in the routing table:

```
hostname# show route
        ....
        R   192.168.32.0/24 [120/4] via 10.1.1.2
        O   192.168.32.0/19 [110/229840] via 10.1.1.3
        ....
```

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but the 192.168.32.0/24 has the longest prefix within the routing table (24 bits verses 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.

Dynamic Routing and Failover

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update

messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a router to which all unroutable packets are sent), for example, can be designated to act as a repository for all unroutable packets, ensuring that all messages are at least handled in some way.

Dynamic routes are synchronized on the standby unit when the routing table changes on the active unit, which means that all additions, deletions, or changes on the active unit are immediately propagated to the standby unit. If the standby unit becomes active after the primary unit has been active for a period of time, routes become synchronized as a part of the failover bulk synchronization process, so the routing table on the active/standby failover pair should appear the same.

For more information about static routes and how to configure them, see the "Configuring Static and Default Routes" section on page 31-2.

Dynamic Routing and Clustering

Dynamic routing is fully integrated in a cluster, and routes are shared across units (up to eight units are allowed in a cluster). Routing table entries are also replicated across units in a cluster.

When one unit transitions from the slave to the master, the epoch number (32-bit sequence number) for the RIB table is incremented. After the transition, the new master unit initially has RIB table entries that are the mirror image of the previous master unit. In addition, the reconvergence timer starts on the new master unit. When the epoch number for the RIB table is incremented, all existing entries are considered stale. Forwarding of IP packets continues as normal. On the new master unit, dynamic routing protocols start to either update existing route entries or create new route entries with the new epoch number. These modified or new entries with the current epoch number indicate that they have been refreshed and are synchronized to all slave units. After the reconvergence timer has expired, old entries from the RIB table are removed. RIB table entries for OSPF routes, RIP routes, and EIGRP routes are synchronized to the slave units.

Bulk synchronization occurs only when a unit joins a cluster and is from the master unit to a joining unit.

For dynamic routing updates, when the master unit learns a new route through OSPF, RIP or EIGRP, the master unit sends those updates to all slave units through reliable message transmission. Slave units update their RIB tables after they receive a cluster route update message.

For the supported dynamic routing protocols (OSPF, RIP, and EIGRP), routing packets from layer 2 load balancing interfaces on the slave units are forwarded to the master unit. Only the master unit sees and processes dynamic routing protocol packets. When the slave unit requests a bulk synchronization, all routing entries learned through layer 2 load balancing interfaces are replicated.

When new routing entries are learned through layer 2 load balancing interfaces on the master unit, the new entries are broadcast to all slave units. When existing routing entries are modified because of a network topology change, the modified entries are also synchronized to all slave units. When existing routing entries are removed because of a network topology change, the removed entries are also synchronized to all slave units.

When a combination of layer 2 and layer 3 load balancing interfaces are deployed and configured for dynamic routing, the slave units only have partial topology and neighbor information (including details that were obtained through layer 3 load balancing interfaces) in the routing process because only RIB table entries are synchronized from the master unit for layer 2 load balancing interfaces. You must configure the network to have layer 2 and layer 3 belong to different routing processes and redistribute the load from each routing process.

Table 30-2 provides a summary of the supported configurations. Yes indicates that the combination of two processes (one process to layer 2 and one process to layer 3) work, No indicates that the combination of two processes does not work.

Table 30-2 Summary of Supported Configurations

Layer 2 or Layer 3	OSPF (Layer 3)	EIGRP (Layer 3)	RIP (Layer 3)
OSPF (layer 2)	Yes	Yes	Yes
EIGRP (layer 2)	Yes	No	Yes
RIP (layer 2)	Yes	Yes	No

All the units in a cluster must be in the same mode: either single or multiple context mode. In multiple context mode, the master-slave synchronization includes all the contexts and the RIB table entries of all the contexts in the synchronization message.

In clustering, if you have configured a layer 3 interface, you must also configure the router-id pool setting.

For more information about dynamic routing and clustering, see Chapter 10, "Configuring a Cluster of ASAs."

Dynamic Routing in Multiple Context Mode

In multiple context mode, each context maintains a separate routing table and routing protocol databases. This enables you to configure OSPFv2 and EIGRP independently in each context. You can configure EIGRP in some contexts and OSPFv2 in the same or different contexts. In mixed context mode, you can enable any of the dynamic routing protocols in contexts that are in routed mode. RIP and OSPFv3 are not supported in multiple context mode.

The following table lists the attributes for EIGRP, OSPFv2, route maps used for distributing routes into OSPFv2 and EIGRP processes, and prefix lists used in OSPFv2 to filter the routing updates entering or leaving an area when they are used in multiple context mode:

EIGRP	OSPFv2	Route Maps and Prefix Lists
One instance is supported per context.	Two instances are supported per context.	N/A
It is disabled in the system context.		N/A
Two contexts may use the same or different autonomous system numbers.	Two contexts may use the same or different area IDs.	N/A
Shared interfaces in two contexts may have multiple EIGRP instances running on them.	Shared interfaces in two contexts may have multiple OSPF instances running on them.	N/A
The interaction of EIGRP instances across shared interfaces is supported.	The interaction of OSPFv2 instances across shared interfaces is supported.	N/A

EIGRP (continued)	OSPFv2 (continued)	Route Maps and Prefix Lists (continued)
All CLIs that are available in single mode are also available in multiple context mode.		
Each CLI has an effect only in the context in which it is used.		

Route Resource Management

A resource class called *routes* has been introduced, which specifies the maximum number of routing table entries that can exist in a context. This resolves the problem of one context affecting the available routing table entries in another context and also allows you greater control over the maximum route entries per context.

Because there is no definitive system limit, you can only specify an absolute value for this resource limit; you may not use a percentage limit. Also, there are no minimum and maximum limits per context, so the default class does not change. If you add a new route for any of the static or dynamic routing protocols (connected, static, OSPF, EIGRP, and RIP) in a context and the resource limit for that context is exhausted, then the route addition fails and a syslog message is generated.

Disabling Proxy ARPs

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 used when a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. The ASA uses proxy ARP when you configure NAT and specify a mapped address that is on the same network as the ASA interface. The only way traffic can reach the hosts is if the ASA uses proxy ARP to claim that the MAC address is assigned to destination mapped addresses.

Under rare circumstances, you might want to disable proxy ARP for NAT addresses.

If you have a VPN client address pool that overlaps with an existing network, the ASA by default sends proxy ARPs on all interfaces. If you have another interface that is on the same Layer 2 domain, it will see the ARP requests and will answer with the MAC address of its interface. The result of this is that the return traffic of the VPN clients towards the internal hosts will go to the wrong interface and will get dropped. In this case, you need to disable proxy ARPs for the interface on which you do not want proxy ARPs.

To disable proxy ARPs,perform the following steps:

Step 1 Choose **Configuration > Device Setup > Routing > Proxy ARPs**.

The Interface field lists the interface names. The Proxy ARP Enabled field shows whether or not proxy ARP is enabled (Yes) or disabled (No) for NAT global addresses.

Step 2 To enable proxy ARP for the selected interface, click **Enable**. By default, proxy ARP is enabled for all interfaces.

Step 3 To disable proxy ARP for the selected interface, click **Disable**.

Step 4 Click **Apply** to save your settings to the running configuration.

Configuring Static and Default Routes

This chapter describes how to configure static and default routes on the ASA and includes the following sections:

- Information About Static and Default Routes, page 31-1
- Licensing Requirements for Static and Default Routes, page 31-2
- Guidelines and Limitations, page 31-2
- Configuring Static and Default Routes, page 31-2
- Monitoring a Static or Default Route, page 31-8
- Configuration Examples for Static or Default Routes, page 31-8
- Feature History for Static and Default Routes, page 31-9

Information About Static and Default Routes

To route traffic to a nonconnected host or network, you must define a static route to the host or network or, at a minimum, a default route for any networks to which the ASA is not directly connected; for example, when there is a router between a network and the ASA.

Without a static or default route defined, traffic to nonconnected hosts or networks generates the following syslog message:

```
%ASA-6-110001: No route to dest_address from source_address
```

You might want to use static routes in single context mode in the following cases:

- Your networks use a different router discovery protocol from EIGRP, RIP, or OSPF.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.

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 outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the ASA.

In transparent firewall mode, for traffic that originates on the ASA and is destined for a nondirectly connected network, you need to configure either a default route or static routes so the ASA knows out of which interface to send traffic. Traffic that originates on the ASA might include communications to a

syslog server, Websense or N2H2 server, or AAA server. If you have servers that cannot all be reached through a single default route, then you must configure static routes. Additionally, the ASA supports up to three equal cost routes on the same interface for load balancing.

Licensing Requirements for Static and Default Routes

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.

IPv6 Guidelines

Supports IPv6.

Failover Guidelines

Supports Stateful Failover of dynamic routing protocols.

Additional Guidelines

- IPv6 static routes are not supported in transparent mode in ASDM.
- In clustering, static route monitoring is only supported on the master unit. For information about clustering, see Chapter 10, “Configuring a Cluster of ASAs.”

Configuring Static and Default Routes

This section explains how to configure a static route and a static default route and includes the following topics:

- Configuring a Static Route, page 31-3
- Configuring a Default Static Route, page 31-6
- Configuring IPv6 Default and Static Routes, page 31-7

Configuring a Static Route

Static routing algorithms are basically table mappings established by the network administrator before the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple. Because of this fact, static routing systems cannot react to network changes.

Static routes remain in the routing table even if the specified gateway becomes unavailable. 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 specified interface goes down, and are reinstated when the interface comes back up.

Note If you create a static route with an administrative distance greater than the administrative distance of the routing protocol running on the ASA, 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.

You can define up to three equal cost routes to the same destination per interface. Equal-cost multi-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.

To configure a static route, choose one of the following:

- Adding or Editing a Static Route, page 31-3
- Configuring Static Route Tracking, page 31-5
- Deleting Static Routes, page 31-6

Adding or Editing a Static Route

To add or edit a static route in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Static Routes**.

Step 2 Choose which route to filter by clicking one of the following radio buttons:

- **Both** (filters both IPv4 and IPv6)
- **IPv4 only**
- **IPv6 only**

By default, the Both radio button is selected, 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 **Add** or **Edit**.

The Add or Edit Static Route dialog box appears.

Step 4 From the Interface drop-down list, choose the internal or external network interface name enabled in the Interface field:

- **management** (internal interface)
- **outside** (external interface)

Step 5 In the IP Address field, type an internal or external network IP address for the destination network.

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 In the Gateway IP field, 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 ASA interface is used as the gateway IP address, the ASA will ARP the designated IP address in the packet instead of ARPing the gateway IP address.

The addresses you specify for the static route are the addresses that are in the packet before entering the ASA and performing NAT.

Step 7 Choose the netmask from the drop-down list for the destination network. 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. Enter **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 In the Metric field, type the metric, or administrative distance.

The metric or distance is the administrative distance for the route. The default is 1 if you do not specify a value. Administrative distance is a parameter used to compare routes among different routing protocols. The default administrative distance for static routes is 1, giving it precedence over routes discovered by dynamic routing protocols, but not directly connected routes.

The default administrative distance for routes discovered by OSPF is 110. If a static route has the same administrative distance as a dynamic route, the static routes take precedence. Connected routes always take precedence over static or dynamically discovered routes.

Step 9 (Optional) In the Options area, choose one of the following options for a static route:

- **None** to have no options specified for the static route. This setting is the default.
- **Tunneled** to specify the route as the default tunnel gateway for VPN traffic. This setting is used for the default route only. You can configure only one tunneled route per device. The tunneled option is not supported in transparent mode.
- **Tracked** to specify that the route is tracked. The tracking object ID and the address of the tracking target also appear. The tracked option is supported in single, routed mode only. Specify the following settings for the tracked option:
  - In the Track ID field, enter a unique identifier for the route tracking process.
  - In the Track IP Address/DNS Name field, 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 from that interface.
  - In the SLA ID field, enter a unique identifier for the SLA monitoring process.

Note The Tracked option is not supported for IPv6.

Step 10 (Optional) Click **Monitoring Options**.

The Route Monitoring Options dialog box appears. From here, you change the following tracking object monitoring properties:

- Frequency, which allows you to modify how often, in seconds, the ASA should test for the presence of the tracking target. Valid values range from 1 to 604800 seconds. The default value is 60 seconds.
- Threshold, which allows you to enter the amount of time, in milliseconds, that indicates an over-threshold event. This value cannot be more than the timeout value.
- Timeout, which allows you to modify the amount of time, in milliseconds, that the route monitoring operation should wait for a response from the request packets. Valid values range from 0 to 604800000 milliseconds. The default value is 5000 milliseconds.
- Data Size, which allows you to modify the size of data payload to use in the echo request packets. The default value is 28. Valid values range 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, which allows you to choose a value for the type of service byte in the IP header of the echo request. Valid values are from 0 to 255. The default value is 0.
- Number of Packets, which allows you to choose the number of echo requests to send for each test. Valid values range from 1 to 100. The default value is 1.

Step 11 Click **OK**.

Step 12 Click **Apply** to save the configuration.

The added or edited route information appears in the Static Routes pane. The monitoring process begins as soon as you save the newly configured route.

Configuring Static Route Tracking

To configure tracking for a static route, perform the following steps:

Note Static route tracking is available for IPv4 routes only.

Step 1 Choose a target of interest. Make sure that the target responds to echo requests.

Step 2 Open the Static Routes pane by choosing **Configuration > Device Setup > 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 settings for this route.

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.

Step 6 (Optional) To configure the monitoring properties, click **Monitoring Options** in the Add Static Route dialog box.

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 by repeating Steps 1 through 7.

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.

Step 9 Click **OK** to save your changes.

Deleting Static Routes

To delete a static route, perform the following steps:

Step 1 Choose **Configuration > Device Setup > Routing > Static Routes**.

Step 2 On the Static Routes pane, choose 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**.

The deleted route is removed from list of routes on in the main Static Routes pane.

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

Configuring a Default Static Route

A default route identifies the gateway IP address to which the ASA sends all IP packets for which it does not have a learned or static route. A default static route is simply a static route with 0.0.0.0/0 as the destination IP address. Routes that identify a specific destination take precedence over the default route.

Note In Versions 7.0(1) and later, if you have two default routes configured on different interfaces that have different metrics, the connection to the ASA that is made from the higher metric interface fails, but connections to the ASA from the lower metric interface succeed as expected.

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 a default route with a different interface than a previously defined default route, you receive the following message:

```
"ERROR: Cannot add route entry, possible conflict with existing routes."
```

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 traffic from a tunnel terminating on the ASA that cannot be routed using learned or static routes is sent to this route. For traffic emerging from a tunnel, this route overrides any other configured or learned default routes.

Limitations on Configuring a Default Static Route

The following restrictions apply to default routes with the tunneled option:

- Do not enable unicast RPF (**ip verify reverse-path** command) on the egress interface of a tunneled route, because this setting causes the session to fail.
- Do not enable TCP intercept on the egress interface of the tunneled route, because this setting causes the session to fail.
- Do not use the VoIP inspection engines (CTIQBE, H.323, GTP, MGCP, RTSP, SIP, SKINNY), the DNS inspect engine, or the DCE RPC inspection engine with tunneled routes, because these inspection engines ignore the tunneled route.
- You cannot define more than one default route with the tunneled option.
- ECMP for tunneled traffic is not supported.

To add or edit a tunneled default static route in ASDM, perform the following steps:

Step 1 On the main ASDM window, choose **Configuration > Device Setup > Routing > Static Routes**.

Step 2 Click **Add** or **Edit**.

Step 3 In the Options area, choose **Tunneled**.

Step 4 Click **OK**.

Configuring IPv6 Default and Static Routes

The ASA automatically routes IPv6 traffic between directly connected hosts if the interfaces to which the hosts are attached are enabled for IPv6 and the IPv6 ACLs allow the traffic.

To add or edit a default static route in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Static Routes**.

Step 2 Click the **IPv6 only** radio button.

Step 3 Click **Add** or **Edit**.

Step 4 Click **OK**.

Monitoring a Static or Default Route

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway becomes unavailable. Static routes are only removed from the routing table if the associated interface on the ASA 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. For example, you can 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 ASA implements this feature by associating a static route with a monitoring target that you define, and monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any network object that you choose, but you should consider using the following:

- 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 on the target network, such as a AAA server, that the ASA needs to communicate with
- A persistent network object on the destination network

Note A desktop or notebook computer that may be shut down at night is not a good choice.

You can configure static route tracking for statically defined routes or default routes obtained through DHCP or PPPoE. You can only enable PPPoE clients on multiple interfaces with route tracking configured.

To monitor the state of a route in ASDM, in the main ASDM window, perform the following steps:

Step 1 Choose **Monitoring > Routing > Routes**.

In the Routes pane, each row represents one route. You can filter by IPv4 connections, IPv6 connections, or both. The routing information includes the protocol, the route type, the destination IP address, the netmask or prefix length, the gateway IP address, the interface through which the route is connected, and the administrative distance.

Step 2 To update the current list, click **Refresh**.

Configuration Examples for Static or Default Routes

The following example shows how to create a static route that sends all traffic destined for 10.1.1.0/24 to the router 10.1.2.45, which is connected to the inside interface, defines three equal cost static routes that direct traffic to three different gateways on the outside interface, and adds a default route for tunneled traffic. The ASA then distributes the traffic among the specified gateways:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Static Routes**.

Step 2 Choose **Management** from the Interfaces drop-down list.

Step 3 Enter **10.1.1.0** in the IP Address field.

Step 4 Choose **255.255.255.0** from the Mask drop-down list.

Step 5 Enter **10.1.2.45 1** in the Gateway IP field.

A static route is created that sends all traffic destined for 10.1.1.0/24 to the router 10.1.2.45, which is connected to the inside interface.

Step 6 Click **OK**.

Step 7 Choose **Configuration > Device Setup > Routing > Static Routes**.

Step 8 Click **Add**.

Step 9 Enter the IP Address in the IP Address field for the destination network.

In this case, the route IP addresses are: 192.168.2.1, 192.168.2.2, 192.168.2.3, and 192.168.2.4. When adding 192.168.2.4, click the **Tunneled** radio button in the Options area.

Step 10 Enter the Gateway IP Address in the Gateway IP Address field for the address of the next hop router.

The addresses you specify for the static route are the addresses that are in the packet before entering the ASA and performing NAT.

Step 11 Choose the netmask for the destination network from the NetMask drop-down list.

Step 12 Click **OK**.

Feature History for Static and Default Routes

Table 31-1 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 31-1 Feature History for Static and Default Routes

Feature Name	Platform Releases	Feature Information
Routing	7.0(1)	Static and default routing were introduced. We introduced the following screen: Configuration > Device Setup > Routing.
Clustering	9.0(1)	Supports static route monitoring on the master unit only.

Defining Route Maps

This chapter describes route maps and includes the following sections:

- Information About Route Maps, page 32-1
- Licensing Requirements for Route Maps, page 32-3
- Guidelines and Limitations, page 32-3
- Defining a Route Map, page 32-4
- Customizing a Route Map, page 32-5
- Configuration Example for Route Maps, page 32-8
- Feature History for Route Maps, page 32-8

Information About Route Maps

Route maps are used when redistributing routes into an OSPF, RIP, or EIGRP routing process. They are also used when generating a default route into an OSPF routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, each has a permit or deny result. Evaluation of ACL or route maps consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms—Criteria matches and match interpretation are dictated by the way that they are applied. The same route map applied to different tasks might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps frequently use ACLs as matching criteria.
- The main result from the evaluation of an access list is a yes or no answer—An ACL either permits or denies input data. Applied to redistribution, an ACL determines if a particular route can (route matches ACLs permit statement) or can not (matches deny statement) be redistributed. Typical route maps not only permit (some) redistributed routes but also modify information associated with the route, when it is redistributed into another protocol.
- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.

- Each ACL ends with an implicit deny statement, by design convention; there is no similar convention for route maps. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Fortunately, route maps that are applied to redistribution behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained deny statement at the end.

The dynamic protocol **redistribute** command allows you to apply a route map. In ASDM, this capability for redistribution can be found when you add or edit a new route map (see the "Defining a Route Map" section on page 32-4). Route maps are preferred if you intend to either modify route information during redistribution or if you need more powerful matching capability than an ACL can provide. If you simply need to selectively permit some routes based on their prefix or mask, we recommends that you use a route map to map to an ACL (or equivalent prefix list) directly in the **redistribute** command. If you use a route map to selectively permit some routes based on their prefix or mask, you typically use more configuration commands to achieve the same goal.

Note You must use a standard ACL as the match criterion for your route map. Using an extended ACL will not work, and your routes will never be redistributed. We recommend that you number clauses in intervals of 10, to reserve numbering space in case you need to insert clauses in the future.

This section includes the following topics:

- Permit and Deny Clauses, page 32-2
- Match and Set Clause Values, page 32-2

Permit and Deny Clauses

Route maps can have permit and deny clauses. In the **route-map ospf-to-eigrp** command, there is one deny clause (with sequence number 10) and two permit clauses. The deny clause rejects route matches from redistribution. Therefore, the following rules apply:

- If you use an ACL in a route map using a permit clause, routes that are permitted by the ACL are redistributed.
- If you use an ACL in a route map deny clause, routes that are permitted by the ACL are not redistributed.
- If you use an ACL in a route map permit or deny clause, and the ACL denies a route, then the route map clause match is not found and the next route-map clause is evaluated.

Match and Set Clause Values

Each route map clause has two types of values:

- A match value **s**elects routes to which this clause should be applied.
- A set value modifies information that will be redistributed into the target protocol.

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeed, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes might be modified by the values set from the Set Value tab in ASDM or from the **set** commands. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found whose **match** command(s), or Match Clause as set from the Match Clause tab in ASDM, match the route or until the end of the route map is reached.

A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several **match** commands or Match Clause values in ASDM are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a **match** command or Match Clause value in ASDM refers to several objects in one command, either of them should match (the logical OR algorithm is applied). For example, in the **match ip address 101 121** command, a route is permitted if access list 101 or access list 121 permits it.
- If a **match** command or Match Clause value in ASDM is not present, all routes match the clause. In the previous example, all routes that reach clause 30 match; therefore, the end of the route map is never reached.
- If a **set** command, or Set Value in ASDM, is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.

Note Do not configure a **set** command in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a **match** or **set** command, or Match or Set Value as set on the Match or Set Value tab in ASDM, performs an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allows a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).

Licensing Requirements for Route Maps

The following table shows the licensing requirements for route maps:

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 context mode and multiple context mode.

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

Route maps do not support access lists that include a user, user group, or fully qualified domain name objects.

Defining a Route Map

You must define a route map when specifying which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Adding or Editing a Route Map

In ASDM, you can define a route map by adding, editing, or deleting a route map name, sequence number, or redistribution.

To add or edit a route map, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

The Add Route Map or Edit Route Map dialog box appears.

Step 3 Enter the route map name and sequence number. The route map name is the name that you assign to a particular route. The sequence number is the order in which you add or delete the route map entries into the ASA.

Note If you are editing an existing route map, the fields for Route Map name and sequence number are already filled in.

Step 4 To reject route matches from redistribution, click **Deny**. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. To allow route matches for redistribution. click **Permit**. If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed.

In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.

Step 5 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
  - Enter the interface name in the Interface field, or click the ellipses to display the Browse Interface dialog box.
  - Choose the interface type (**inside** or **outside**), click **Selected Interface**, then click **OK**.
  - Check the **Match IP Address** check box to enable or disable the Match address of a route or match packet.
  - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.

  - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
  - Choose Access List to Prefix List from the drop-down list to match the IP address.
  - According to the previous selection, click the ellipses to display the Browse Access List or Browse Prefix List dialog box.
  - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
  - In the Metric Value field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.
- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Step 6 Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:

- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and type the value in the Value field.
- Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.

Customizing a Route Map

This section describes how to customize the route map and includes the following topics:

- Defining a Route to Match a Specific Destination Address, page 32-5
- Configuring Prefix Lists, page 32-6
- Configuring Prefix Rules, page 32-7
- Configuring the Metric Values for a Route Action, page 32-7

Defining a Route to Match a Specific Destination Address

To define a route to match a specified destination address, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

The Add Route Map dialog box appears. From this dialog box you can assign or choose the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add the entries.

Step 3 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
  - Enter the interface name in the Interface field, or click the ellipses to display the Browse Interface dialog box.
  - Choose the interface type (**inside** or **outside**), click **Selected Interface**, then click **OK**.
  - Check the **Match IP Address** check box to enable or disable the Match address of a route or match packet.
  - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.
  - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
  - Choose Access List to Prefix List from the drop-down list to match the IP address.
  - According to the previous selection, click the ellipses to display the Browse Access List or Browse Prefix List dialog box.
  - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
  - In the Metric Value field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.
- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Configuring Prefix Lists

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

To add prefix lists, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Prefix Rules**.

Step 2 From the Add drop-down list, choose **Add Prefix List**.

The Add Prefix List dialog box appears.

Step 3 Enter the prefix name and description, then click **OK**.

Configuring Prefix Rules

Note You must configure a prefix list before you may configure a prefix rule.

To configure prefix rules, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Prefix Rules**.

Step 2 Click **Add**.

The Add Prefix Rule dialog box appears. From this dialog box, you can add a sequence number, specify a prefix for the network, its redistribution access (that is, permit or deny), the minimum and maximum prefix length.

Step 3 Enter an optional sequence number or accept the default value.

Step 4 Specify the prefix number in the format of IP address/mask length.

Step 5 Click the **Permit** or **Deny** radio button to indicate the redistribution access.

Step 6 Enter the optional minimum and maximum prefix lengths.

Step 7 Click **OK** when you are done.

The new or revised prefix rule appears in the list.

Step 8 Check the **Enable Prefix list sequence numbering** check box if you want to use automatically generated sequence numbers.

Step 9 Click **Apply** to save your changes.

Configuring the Metric Values for a Route Action

To configure the metric value for a route action, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

The Add Route Map or Edit Route Map dialog box appears. From this dialog box, you can assign or select the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add route map entries.

Step 3 Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:

- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and enter the value in the Value field.
- Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.

Configuration Example for Route Maps

The following example shows how to redistribute routes with a hop count equal to 1 into OSPF.

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

Step 3 Enter **1-to-2** in the Route Map Name field.

Step 4 Enter the routing sequence number in the Sequence Number field.

Step 5 Click the **Permit** radio button.

By default this tab is on top.

Step 6 Click the **Match Clause** tab.

Step 7 Check the **Match Metric of Route** check box and type **1** for the metric value.

Step 8 Click the **Set Clause** tab.

Step 9 Check the **Set Metric Value** check box, and type **5** for the metric value.

Step 10 Check the **Set Metric-Type** check box, and choose **Type-1**.

Feature History for Route Maps

Table 26-1 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 32-1 Feature History for Route Maps

Feature Name	Platform Releases	Feature Information
Route maps	7.0(1)	We introduced this feature. We introduced the following screen: Configuration > Device Setup > Routing > Route Maps.
Enhanced support for static and dynamic route maps	8.0(2)	Enhanced support for dynamic and static route maps was added.
Dynamic Routing in Multiple Context Mode	9.0(1)	Route maps are supported in multiple context mode.

Configuring OSPF

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

The chapter includes the following sections:

- Information About OSPF, page 33-1
- Licensing Requirements for OSPF, page 33-3
- Guidelines and Limitations, page 33-3
- Configuring OSPFv2, page 33-5
- Customizing OSPFv2, page 33-6
- Configuring OSPFv3, page 33-21
- Removing the OSPF Configuration, page 33-32
- Configuration Example for OSPFv2, page 33-32
- Configuration Example for OSPFv3, page 33-34
- Monitoring OSPF, page 33-35
- Additional References, page 33-36
- Feature History for OSPF, page 33-37

Information About 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 uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

The advantages of OSPF over RIP include the following:

- OSPF link-state database updates are sent less frequently than RIP updates, and the link-state database is updated instantly, rather than gradually, as stale information is timed out.

- Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The ASA calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.

The disadvantage of shortest path first algorithms is that they require a lot of CPU cycles and memory.

The ASA can run two processes of OSPF protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to coexist, but OSPF does not allow overlapping addresses). Or you might want to run one process on the inside and another on the outside, and redistribute a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPF routing process from another OSPF routing process, a RIP routing process, or from static and connected routes configured on OSPF-enabled interfaces.

The ASA supports the following OSPF features:

- Intra-area, interarea, and external (Type I and Type II) routes.
- Virtual links.
- LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Configuring the ASA as a designated router or a designated backup router. The ASA also can be set up as an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

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 (such as 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 ASA acting as an ABR. Type 3 LSAs (interarea routes) can be filtered from one area to other, which allows you to use NAT and OSPF together without advertising private networks.

Note Only Type 3 LSAs can be filtered. If you configure the ASA 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 ASA. Also, you should not mix public and private networks on the same ASA interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the ASA at the same time.

Implementation Differences Between OSPFv2 and OSPFv3

OSPFv3 is not backward compatible with OSPFv2. To use OSPF to route both IPv4 and IPv6 traffic, you must run both OSPFv2 and OSPFv3 at the same time. They coexist with each other, but do not interact with each other.

The additional features that OSPFv3 provides include the following:

- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Using Clustering

For more information about dynamic routing and clustering, see the "Dynamic Routing and Clustering" section on page 30-9.

For more information about using clustering, see Chapter 10, "Configuring a Cluster of ASAs."

Licensing Requirements for OSPF

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

OSPFv2 supports single and multiple context mode.

OSPFv3 supports single mode only.

Firewall Mode Guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

Failover Guidelines

OSPFv3 supports Stateful Failover.

IPv6 Guidelines

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The ASA installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.
- OSPFv3 packets can be filtered out using IPv6 ACLs in the **capture** command.

Clustering Guidelines

- OSPFv2 and OSPFv3 support clustering.
- OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment.
- In the spanned interface mode, dynamic routing is not supported on management-only interfaces.
- In individual interface mode, make sure that you establish the master and slave units as either OSPFv2 or OSPFv3 neighbors.
- When you configure both OSPFv2 and EIGRP, you can use either spanned interface mode or individual interface mode; you cannot use the two modes at the same time.
- In individual interface mode, OSPFv2 adjacencies can only be established between two contexts on a shared interface on the master unit. Configuring static neighbors is supported only on point-to point-links; therefore, only one neighbor statement is allowed on an interface.
- The router ID is optional in the OSPFv2, OSPFv3, and EIGRP router configuration mode. If you do not explicitly set a router ID, then a router ID is automatically generated and set to the highest IPv4 address on any data interface in each of the cluster units.
- If the cluster interface mode has not been configured, then only a single, dotted-decimal IPv4 address is allowed as the router ID, and the **cluster pool** option is disabled.
- If the cluster interface mode is set to a spanned configuration, then only a single, dotted-decimal IPv4 address is allowed as the router ID, and the **cluster pool** option is disabled.
- If the cluster interface mode is set to an individual configuration, then the **cluster pool** option is mandatory, and a single, dotted-decimal IPv4 address is not allowed as the router ID.
- When the cluster interface mode is changed from a spanned to an individual configuration and vice versa without specifying the **check-detail** or **nocheck** options, then the entire configuration including the router ID is removed.
- If any of the dynamic routing protocol router ID configurations are incompatible with the new interface mode, then an error message appears on the console and the interface mode CLI fails. The error message has one line per dynamic routing protocol (OSPFv2, OSPFv3, and EIGRP) and lists the names of each context in which the incompatible configuration occurs.
- If the **nocheck** option is specified for the **cluster interface mode** command, then the interface mode is allowed to change although all the router ID configurations may not be compatible with the new mode.
- When the cluster is enabled, the router ID compatibility checks are repeated. If any incompatibility is detected, then the **cluster enable** command fails. The administrator needs to correct the incompatible router ID configuration before the cluster can be enabled.

- When a unit enters a cluster as a slave, then we recommend that you specify the **nocheck** option for the **cluster interface mode** command to avoid any router ID compatibility check failures. The slave unit still inherits the router configuration from the master unit.
- When a mastership role change occurs in the cluster, the following behavior occurs:
  - In spanned interface mode, the router process is active only on the master unit and is in a suspended state on the slave units. Each cluster unit has the same router ID because the configuration has been synchronized from the master unit. As a result, a neighboring router does not notice any change in the router ID of the cluster during a role change.
  - In individual interface mode, the router process is active on all the individual cluster units. Each cluster unit chooses its own distinct router ID from the configured cluster pool. A mastership role change in the cluster does not change the routing topology in any way.

Additional Guidelines

- OSPFv2 and OSPFv3 support multiple instances on an interface.
- OSPFv3 supports encryption through ESP headers in a non-clustered environment.
- OSPFv3 supports Non-Payload Encryption.

Configuring OSPFv2

This section describes how to enable an OSPFv2 process on the ASA.

After you enable OSPFv2, you need to define a route map. For more information, see the "Defining a Route Map" section on page 32-4. Then you generate a default route. For more information, see the "Configuring Static and Default Routes" section on page 31-2.

After you have defined a route map for the OSPFv2 process, you can customize it for your particular needs, To learn how to customize the OSPFv2 process on the ASA, see the "Customizing OSPFv2" section on page 33-6.

To enable OSPFv2, you need to create an OSPFv2 routing process, specify the range of IP addresses associated with the routing process, then assign area IDs associated with that range of IP addresses.

You can enable up to two OSPFv2 process instances. Each OSPFv2 process has its own associated areas and networks.

To enable OSPFv2, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

In the OSPF Setup pane, you can enable OSPF processes, configure OSPF areas and networks, and define OSPF route summarization.

Step 2 The three tabs in ASDM used to enable OSPF are as follows:

- The Process Instances tab allows you to enable up to two OSPF process instances for each context. Single context mode and multiple context mode are both supported. After you check the **Enable Each OSPF Process** check box, you can enter a unique identifier numeric identifier for that OSPF process. This process ID is used internally and does not need to match the OSPF process ID on any other OSPF devices; valid values range from 1 to 65535. Each OSPF process has its own associated areas and networks.

If you click **Advanced**, the Edit OSPF Process Advanced Properties dialog box appears. From here, you can configure the Router ID, cluster IP address pools in Layer 2 or Layer 3 clustering, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for each OSPF process.

- The Area/Networks tab allows you to display the areas and the networks that they include for each OSPF process on the ASA. From this tab you can display the area ID, the area type, and the type of authentication set for the area. To add or edit the OSPF area or network, see the "Configuring OSPFv2 Area Parameters" section on page 33-13 for more information.
- The Route Summarization tab allows you to configure an ABR. 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 so that they are contiguous, you can configure the ABR to advertise a summary route that includes all the individual networks within the area that fall into the specified range. See the "Configuring Route Summarization Between OSPFv2 Areas" section on page 33-10 for more information.

Customizing OSPFv2

This section explains how to customize the OSPFv2 processes and includes the following topics:

- Redistributing Routes Into OSPFv2, page 33-6
- Configuring Route Summarization When Redistributing Routes Into OSPFv2, page 33-8
- Configuring Route Summarization Between OSPFv2 Areas, page 33-10
- Configuring OSPFv2 Interface Parameters, page 33-11
- Configuring OSPFv2 Area Parameters, page 33-13
- Configuring an OSPFv2 NSSA, page 33-14
- Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3), page 33-15
- Defining Static OSPFv2 Neighbors, page 33-17
- Configuring Route Calculation Timers, page 33-18
- Logging Neighbors Going Up or Down, page 33-18
- Configuring Filtering in OSPF, page 33-19
- Configuring a Virtual Link in OSPF, page 33-20

Redistributing Routes Into OSPFv2

The ASA can control the redistribution of routes between OSPFv2 routing processes.

Note If you want to redistribute a route by defining which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process, you must first generate a default route. See the "Configuring Static and Default Routes" section on page 31-2, and then define a route map according to the "Defining a Route Map" section on page 32-4.

To redistribute static, connected, RIP, or OSPFv2 routes into an OSPFv2 process, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Redistribution**.

The Redistribution pane displays the rules for redistributing routes from one routing process into an OSPF routing process. 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. You do not need to redistribute static or connected routes if they fall within the range of a network that has been configured through the Setup > Networks tab.

Step 2 Click **Add** or **Edit**.

Alternatively, double-clicking a table entry in the Redistribution pane (if any) opens the Add/Edit OSPF Redistribution Entry dialog box for the selected entry.

Note All steps that follow are optional.

The Add/Edit OSPF Redistribution Entry dialog box lets you add a new redistribution rule 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.

Step 3 Choose the OSPF process associated with the route redistribution entry. If you are editing an existing redistribution rule, you cannot change this setting.

Step 4 Choose the source protocol from which the routes are being redistributed. You can choose one of the following options:

- Static—Redistributes static routes to the OSPF routing process.
- Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS.
- OSPF—Redistributes routes from another OSPF routing process. Choose the OSPF process ID from the list. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or EIGRP routes. Skip to Step 5.
- RIP—Redistributes routes from the RIP routing process.
- EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list.

Step 5 If you have chosen OSPF for the source protocol, choose 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.

Step 6 In the Metric Value field, enter 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.

Step 7 Choose one of the following options for the Metric Type.

- If the metric is a Type 1 external route, choose **1**.
- If the metric is a Type 2 external route, choose **2**.

Step 8 Enter the tag value in the Tag Value field.

The tag value is a 32-bit decimal value attached to each external route that is not used by OSPF itself, but may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.

Step 9 Check the **Use Subnets** check box to enable the redistribution of subnetted routes. Uncheck this check box to cause only routes that are not subnetted to be redistributed.

Step 10 Choose the name of the route map to apply to the redistribution entry from the Route Map drop-down list.

Step 11 If you need to add or configure a route map, click **Manage**.

The Configure Route Map dialog box appears.

Step 12 Click **Add** or **Edit** to define which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. For more information, see the "Defining a Route Map" section on page 32-4.

Step 13 Click **OK**.

Configuring Route Summarization When Redistributing Routes Into OSPFv2

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the ASA to advertise a single route for all the redistributed routes that are included for a specified network address and mask. This configuration decreases the size of the OSPF link-state database.

Routes that match the specified IP address mask pair can be suppressed. The tag value can be used as a match value for controlling redistribution through route maps.

There are two areas that you can configure for route summarization:

- Adding a Route Summary Address, page 33-9
- Adding or Editing an OSPF Summary Address, page 33-9

Adding a Route 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.

Note OSPF does not support summary-address 0.0.0.0 0.0.0.0.

To configure the software advertisement on one summary route for all redistributed routes included for a network address and mask, perform the following steps:

Step 1 In the main ASDM home page, choose **Configuration > Device Setup > Routing > OSPF > Summary Address**.

Step 2 Click **Add**.

The Add OSPF Summary Address Entry dialog box appears. You can add new entries to existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the OSPF Process drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the IP Address field. You cannot change this information when editing an existing entry.

Step 5 Choose the network mask for the summary address from the Netmask drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** 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.

The Tag value displays a 32-bit decimal value that is attached to each external route. This value is not used by OSPF itself, but may be used to communicate information between ASBRs.

Step 7 Click **OK**.

Adding or Editing an OSPF Summary Address

To add or edit OSPF summary address setting, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to 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.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the OSPF Process drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the IP Address field. You cannot change this information when editing an existing entry.

Step 5 Enter the network mask for the summary address from the Netmask drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** 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.

Configuring Route Summarization Between OSPFv2 Areas

Route summarization is the consolidation of advertised addresses. This feature causes a single summary route to be advertised to other areas by an area boundary router. In OSPF, an area boundary router advertises networks in one area into another area. If the network numbers in an area are assigned in a way so that they are contiguous, you can configure the area boundary router to advertise a summary route that includes all the individual networks within the area that fall into the specified range.

To define an address range for route summarization, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to 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.

Step 3 Enter the OSPF Area ID in the Area ID field. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the IP Address field. You cannot change this information when editing an existing entry.

Configuring OSPFv2 Interface Parameters

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the Hello interval, the Dead interval, and the Authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

To configure OSPFv2 interface parameters, perform the following steps:

In ASDM, the Interface pane lets you configure interface-specific OSPF routing properties, such as OSPF message authentication and properties. There are two tabs that help you configure interfaces in OSPF:

- The Authentication tab displays the OSPF authentication information for the ASA interfaces.
- The Properties tab displays the OSPF properties defined for each interface in a table format.

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Interface.**

Step 2 Click the **Authentication** tab to display the authentication information for the ASA interfaces. Double-clicking a row in the table opens the Edit OSPF Authentication Interface dialog box for the selected interface.

Step 3 Click **Edit**.

The Edit OSPF Authentication Interface dialog box appears. The Edit OSPF Interface Authentication dialog box lets you configure the OSPF authentication type and parameters for the selected interface.

Step 4 Choose the Authentication type from the Authentication drop-down list according to the following options:

- **None** to disable OSPF authentication.
- **Authentication Password** to use clear text password authentication (not recommended where security is a concern).
- **MD5** to use MD5 authentication (recommended).
- **Area** (Default) to use the authentication type specified for the area. See the "Configuring OSPFv2 Area Parameters" section on page 33-13 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure this setting.

Step 5 Click the radio button in the Authentication Password area, which includes the settings for entering the password when password authentication is enabled.

a. In the Enter Password field, type a text string of up to eight characters.

b. In the Re-enter Password field, retype the password.

Step 6 Choose the settings for MD5 IDs and keys in the ID area, which includes 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.

a. In the Key ID field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.

b. In the Key field, enter an alphanumeric character string of up to 16 bytes. The key displays for the selected interface.

c. Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 7 Click **OK.**

Step 8 Click the **Properties** tab.

Step 9 Choose the interface that you want to edit. Double-clicking a row in the table opens the Properties tab dialog box for the selected interface.

Step 10 Click **Edit**.

The Edit OSPF Interface Properties dialog box appears. The Interface field displays the name of the interface for which you are configuring OSPF properties. You cannot edit this field.

Step 11 Check or uncheck the **Broadcast** check box to specify that the interface is a broadcast interface.

By default, this check box is checked for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, nonbroadcast interface. Specifying an interface as point-to-point, nonbroadcast 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 the "Defining Static OSPFv2 Neighbors" section on page 33-17 for more information.
- You need to define a static route pointing to the crypto endpoint. See the "Configuring Static and Default Routes" section on page 31-2 for more information.
- If OSPF over a 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 that the OSPF adjacencies can be established over the VPN tunnel.

Step 12 Configure the following options:

- Enter a value in the Cost field, which determines the cost of sending a packet through the interface. The default value is 10.
- In the Priority field, enter the OSPF router priority value.

  When two routers connect to a network, both attempt to become the designated router. The device 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, nonbroadcast interfaces.
- Check or uncheck the **MTU Ignore** check box.

  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.
- Check or uncheck the **Database filter** check box.

  Use this setting to filter the 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 flooding can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents OSPF flooding of the LSA on the selected interface.

Step 13 (Optional) Click **Advanced** to display the Edit OSPF Advanced Interface Properties dialog box, which 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.

Step 14 Enter values for the following:

- The Hello Interval, which specifies the interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected, but 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.
- The Retransmit Interval, which 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.
- The Transmit Delay, which 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.
- The Dead Interval, which 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 that was set in the Hello Interval field.

Configuring OSPFv2 Area Parameters

You can configure several OSPF area parameters. These area parameters (shown in the following task list) include setting authentication, defining stub areas, and assigning specific costs to the default summary route. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area.

To specify OSPFv2 area parameters for your network, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Area/Networks** tab.

The Add OSPF Area dialog box appears.

Step 3 Choose one of the following Area Type options:

- **Normal** to make the area a standard OSPF area. This option is selected by default when you first create an area.

- **Stub** to make 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 (Types 3 and 4) from being flooded into the area by unchecking the Summary check box.
- **Summary** to prevent LSAs from being sent into the stub area when the area being defined is a stub area, uncheck this check box. By default, this check box is checked for stub areas.
- **NSSA** to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the 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 checking the Default Information Originate check box.

Step 4 Enter the IP address in the IP Address field 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 enter **0.0.0.0** in one area.

Step 5 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 6 Choose the OSPF Authentication type from the following options:

- **None** to disable OSPF area authentication. This is the default setting.
- **Password** to provide a clears text password for area authentication, which is not recommended where security is a concern.
- **MD5** to allow MD5 authentication.

Step 7 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 8 Click **OK**.

Configuring an OSPFv2 NSSA

The OSPFv2 implementation of an NSSA is similar to an OSPFv2 stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but it can import autonomous system external routes in a limited way within the area.

NSSA imports Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA ABRs, which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

You can simplify administration if you are an ISP or a network administrator that must connect a central site using OSPFv2 to a remote site that is using a different routing protocol with NSSA.

Before the implementation of NSSA, the connection between the corporate site border router and the remote router could not be run as an OSPFv2 stub area because routes for the remote site could not be redistributed into the stub area, and two routing protocols needed to be maintained. A simple protocol such as RIP was usually run and handled the redistribution. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

Before you use this feature, consider these guidelines:

- You can set a Type 7 default route that can be used to reach external destinations. When configured, the router generates a Type 7 default into the NSSA or the NSSA area boundary router.

- Every router within the same area must agree that the area is NSSA; otherwise, the routers cannot communicate with each other.

To specify area parameters for your network to configure an OSPFv2 NSSA, perform the following steps:

Step 1 From the main ASDM home page, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Area/Networks** tab.

Step 3 Click **Add**.

The Add OSPF Area dialog box appears.

Step 4 Click the **NSSA** radio button in the Area Type area.

Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the 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 checking the Default Information Originate check box.

Step 5 Enter the IP address in the IP Address field 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 enter **0.0.0.0** in one area.

Step 6 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 7 In the Authentication area, click the **None** radio button to disable OSPF area authentication.

Step 8 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 9 Click **OK**.

Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3)

You can assign a range of IPv4 addresses for the router ID cluster pool if you are using Layer 3 clustering.

To assign a range of IPv4 addresses for the router ID cluster pool in Layer 3 clustering for OSPFv2, perform the following steps:

Step 1 From the main ASDM home page, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 Click the **Cluster Pool** radio button. If you are using clustering, then you do not need to specify an IP address pool for the router ID (that is, leave the field blank). If you do not enter an IP address pool, then the ASA uses the automatically generated router ID.

Step 5 Enter the name of the IP address pool, or click the ellipses to display the Select IP Address Pool dialog box.

Step 6 Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

The Add IPv4 Pool dialog box appears.

Step 7 Enter the new IP address pool name in the Name field.

Step 8 Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.

Step 9 Double-click an entry to add it to the Starting IP Address field, then click **OK**.

Step 10 Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box.

Step 11 Double-click an entry to add it to the Ending IP Address field, then click **OK**.

Step 12 Choose the subnet mask from the drop-down list, then click **OK**.

The new IP address pool appears in the Select IP Address Pool list.

Step 13 Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

Step 14 Click **OK**.

Step 15 If you want to change the newly added IP address pool settings, click **Edit**.

The Edit IPv4 Pool dialog box appears.

Step 16 Repeat Steps 4 through 14.

Note You cannot edit or delete an existing IP address pool that has been assigned and is already being used by one or more connection profiles.

Step 17 Click **OK**.

To assign a range of IPv4 addresses for the router ID cluster pool in Layer 3 clustering for OSPFv3, perform the following steps:

Step 1 From the main ASDM home page, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Choose the Cluster Pool option from the Router ID drop-down list. If you do not need to specify an IP address pool for the router ID, choose the Automatic option. If you do not configure an IP address pool, then the ASA uses the automatically generated router ID.

Step 5 Enter the IP address pool name. Alternatively, click the ellipses to display the Select IP Address Pool dialog box.

Step 6 Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

The Add IPv4 Pool dialog box appears.

Step 7 Enter the new IP address pool name in the Name field.

Step 8 Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.

Step 9 Double-click an entry to add it to the Starting IP Address field, then click **OK**.

Step 10 Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box.

Step 11 Double-click an entry to add it to the Ending IP Address field, then click **OK**.

Step 12 Choose the subnet mask from the drop-down list, then click **OK**.

The new IP address pool appears in the Select IP Address Pool list.

Step 13 Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

Step 14 Click **OK**.

Step 15 If you want to change the newly added cluster pool settings, click **Edit**.

The Edit IPv4 Pool dialog box appears.

Step 16 Repeat Steps 4 through 14.

Note You cannot edit or delete an existing IP address pool that has been assigned and is already being used by another OSPFv3 process.

Step 17 Click **OK**.

Defining Static OSPFv2 Neighbors

You need to define static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv2 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Before you begin, you must create a static route to the OSPFv2 neighbor. See Chapter 31, "Configuring Static and Default Routes," for more information about creating static routes.

To define a static OSPFv2 neighbor, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add/Edit OSPF Neighbor Entry dialog box appears. This 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, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPF processes.
- You need to define a static route for each static neighbor.

Step 3 From the OSPF Process drop-down list, choose the OSPF process associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the Neighbor field, enter the IP address of the static neighbor.

Step 5 In the Interface field, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 6 Click **OK**.

Configuring Route Calculation Timers

You can configure the delay time between when OSPFv2 receives a topology change and when it starts an SPF calculation. You also can configure the hold time between two consecutive SPF calculations.

To configure route calculation timers, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA pacing and SPF calculation timers. In the Timers area, enter the following values:

- The SPF Delay Time, which 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.
- The SPF Hold Time, which specifies the hold time between consecutive SPF calculations.Valid values range from 1 to 65534. The default value is 10.
- The LSA Group Pacing, which specifies the interval at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.

Step 5 Click **OK**.

Logging Neighbors Going Up or Down

By default, a syslog message is generated when an OSPFv2 neighbor goes up or down.

To log OSPFv2 neighbors going up or down, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Adjacency Changes area includes settings that define the adjacency changes that cause syslog messages to be sent. In the Adjacency Changes area, enter the following values:

- Check the **Log Adjacency Changes** check box to cause the ASA to send a syslog message whenever an OSPFv2 neighbor goes up or down. This setting is checked by default.
- Check the **Log Adjacency Changes Detail** check box to cause the ASA to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

Step 5 Click **OK**.

Note Logging must be enabled for the neighbor up or down messages to be sent.

Configuring Filtering in OSPF

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 restrict 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.

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Note Only Type 3 LSAs that originate from an ABR are filtered.

To configure filtering in OSPF, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Filtering**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Filtering Entry dialog box lets you add new filters to the Filter table or modify an existing filter. Some of the filtering information cannot be changed when you edit an existing filter.

Step 3 Choose the OSPF process that is associated with the filter entry from the OSPF Process drop-down list.

Step 4 Choose the Area ID that is associated with the filter entry from the Area ID drop-down list. If you are editing an existing filter entry, you cannot modify this setting.

Step 5 Choose a prefix list from the Prefix List drop-down list.

Step 6 Choose the traffic direction being filtered from the Traffic Direction drop-down list.

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.

Step 7 Click **Manage** to display the Configure Prefix Lists dialog box, from which you can add, edit, or delete prefix lists and prefix rules. For more information, see the "Configuring Prefix Lists" section on page 32-6 and the "Configuring Prefix Rules" section on page 32-7.

Step 8 Click **OK**.

Configuring a Virtual Link in OSPF

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.

To define new virtual links or change the properties of existing virtual links, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Virtual Link**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Virtual Link dialog box appears, which allows you to define new virtual links or change the properties of existing virtual links.

Step 3 Choose the OSPF process ID that is associated with the virtual link from the OSPF Process drop-down list. If you are editing an existing virtual link entry, you cannot modify this setting.

Step 4 Choose the Area ID that is associated with the virtual link from the Area ID drop-down list.

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 entry, you cannot modify this setting.

Step 5 In the Peer Router ID field, enter the router ID of the virtual link neighbor.

If you are editing an existing virtual link entry, you cannot modify this setting.

Step 6 Click **Advanced** to edit advanced virtual link properties,

The Advanced OSPF Virtual Link Properties dialog box appears. You can configure the OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings.

Step 7 In the Authentication area, choose the Authentication type by clicking the radio button next to one of the following options:

- **None** to disable OSPF authentication.
- **Authentication Password** to use clear text password authentication. This is not recommended where security is a concern.
- **MD5** to use MD5 authentication (recommended).
- **Area** (Default) to use the authentication type specified for the area. See the "Configuring OSPFv2 Area Parameters" section on page 33-13 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure this setting.

Step 8 In the Authentication Password area, enter and re-enter a password when password authentication is enabled. Passwords must be a text string of up to 8 characters.

Step 9 In the MD5 IDs and Key area, enter 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. Specify the following settings:

a. In the Key ID field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.

b. In the Key field, enter an alphanumeric character string of up to 16 bytes. The Key ID displays for the selected interface.

c. Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 10 In the Interval area, specify the interval timing for the packet by choosing from the following options:

- **Hello Interval** to specify 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** to specify 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** to specify 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** to specify 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.

Step 11 Click **OK**.

Configuring OSPFv3

This section describes how to configure OSPFv3 routing processes and includes the following topics:

- Enabling OSPFv3, page 33-22
- Configuring OSPFv3 Interface Parameters, page 33-22
- Configuring OSPFv3 Area Parameters, page 33-23
- Configuring a Virtual Link Neighbor, page 33-24
- Configuring OSPFv3 Passive Interfaces, page 33-25
- Configuring OSPFv3 Administrative Distance, page 33-26
- Configuring OSPFv3 Timers, page 33-27
- Defining Static OSPFv3 Neighbors, page 33-28
- Sending Syslog Messages, page 33-29
- Suppressing Syslog Messages, page 33-29
- Calculating Summary Route Costs, page 33-29
- Generating a Default External Route into an OSPFv3 Routing Domain, page 33-30
- Configuring an IPv6 Summary Prefix, page 33-30
- Redistributing IPv6 Routes, page 33-31

Enabling OSPFv3

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, then redistribute the route into the targeted OSPFv3 routing processes.

To enable OSPFv3, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 On the Process Instances tab, check the **Enable OSPFv3 Process** check box. You can enable up to two OSPF process instances. Only single context mode is supported.

Step 3 Enter a process ID in the Process ID field. The ID can be any positive integer.

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

Step 5 To continue, see the "Configuring OSPFv3 Area Parameters" section on page 33-23.

Configuring OSPFv3 Interface Parameters

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

To configure OSPFv3 interface parameters for IPv6, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Interfaces**.

Step 2 Click the **Authentication** tab.

Step 3 To specify the authentication parameters for an interface, select the interface and click **Edit**.

The Edit OSPFv3 Interface Authentication dialog box appears.

Step 4 Choose the authentication type from the Authentication Type drop-down list. The available options are Area, Interface, and None. The None option indicates that no authentication is used.

Step 5 Choose the authentication algorithm from the Authentication Algorithm drop-down list. Supported values are SHA-1 and MD5.

Step 6 Enter the authentication key in the Authentication Key field. When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.

Step 7 Choose the encryption algorithm from the Encryption Algorithm drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 8 Enter the encryption key in the Encryption Key field.

Step 9 Click **OK**.

Step 10 Click the **Properties** tab.

Step 11 Select the interface whose properties you want to modify, and click **Edit**.

The Edit OSPFv3 Interface Properties dialog box appears.

Step 12 Check the **Enable OSPFv3 on this interface** check box.

Step 13 Choose the process ID from the drop-down list.

Step 14 Choose the area ID from the drop-down list.

Step 15 (Optional) Specify the area instance ID to be assigned to the interface. An interface can have only one OSPFv3 area. You can use the same area on multiple interfaces, and each interface can use a different area instance ID.

Step 16 Choose the network type from the drop-down list. Supported options are Default, Broadcast, and Point-to-Point.

Step 17 Enter the cost of sending a packet on an interface in the Cost field.

Step 18 Enter the router priority, which helps determine the designated router for a network. in the Priority field. Valid values range from 0 to 255.

Step 19 Check the **Disable MTU mismatch detection** check box to disable the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.

Step 20 Check the **Filter outgoing link state advertisements** check box to filter outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.

Step 21 In the Timers area, in the Dead Interval field, enter the time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range from 1 to 65535.

Step 22 In the Hello Interval field, enter the interval in seconds between hello packets sent on the interface. The value must be the same for all nodes on a specific network and can range from 1 to 65535. The default interval is 10 seconds for Ethernet interfaces and 30 seconds for non-broadcast interfaces.

Step 23 In the Retransmit Interval field, enter the time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

Step 24 In the Transmit Delay field, enter the estimated time in seconds to send a link-state update packet on the interface. Valid values range from 1 to 65535 seconds. The default is 1 second.

Step 25 Click **OK**.

Step 26 Click **Apply** to save your changes.

Configuring OSPFv3 Area Parameters

To configure OSPFv3 area parameters, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Areas** tab.

Step 3 To add a new area, click **Add**. To modify an existing area, click **Edit**. To remove a selected area, click **Delete**.

The Add OSPFv3 Area dialog box or Edit OSPFv3 Area dialog box appears.

Step 4 From the OSPFv3 Process ID drop-down list, choose the process ID.

Step 5 Enter the area ID, which specifies the area for which routes are to be summarized, in the Area ID field.

Step 6 Choose the area type from the Area Type drop-down list. Available options are Normal, NSSA, and Stub.

Step 7 To allow the sending of summary LSAs into the area, check the **Allow sending of summary LSAs into the area** check box.

Step 8 To allow redistribution to import routes to normal and not so stubby areas, check the **Redistribution imports routes to normal and NSSA areas** check box.

Step 9 To generate a default external route into an OSPFv3 routing domain, check the **Default information originate** check box.

Step 10 Enter the metric used for generating the default route in the Metric field. The default value is 10. Valid metric values range from 0 to 16777214.

Step 11 Choose the metric type from the Metric Type drop-down list. The metric type is the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.

Step 12 Enter the cost in the Default Cost field.

Step 13 Click **OK**.

Step 14 Click the **Route Summarization** tab.

Step 15 To specify a new range for consolidating and summarizing routes, click **Add**. To modify an existing range for consolidating and summarizing routes, click **Edit**.

The Add Route Summarization dialog box or Edit Route Summarization dialog box appears.

Step 16 Choose the process ID from the Process ID drop-down list.

Step 17 Choose the area ID from the Area ID drop-down list.

Step 18 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 19 (Optional) Enter the metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.

Step 20 Check the **Advertised** check box to set the address range status to advertised and generate a Type 3 summary LSA.

Step 21 Click **OK**.

Step 22 To continue, see the "Configuring a Virtual Link Neighbor" section on page 33-24.

Configuring a Virtual Link Neighbor

To configure a virtual link neighbor, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Virtual Link**.

Step 2 To add a new virtual link neighbor, click **Add**. To modify an existing virtual link neighbor, click **Edit**. To remove a selected virtual link neighbor, click **Delete**.

The Add Virtual Link dialog box or Edit Virtual Link dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the area ID from the Area ID drop-down list.

Step 5 Enter the peer router ID (that is, the IP address) in the Peer Router ID field.

Step 6 (Optional) Enter the time-to-live (TTL) security hop count on a virtual link in the TTL Security field. The hop count value can range from 1 to 254.

Step 7 In the Timers area, enter the time in seconds that hello packets are not seen before a neighbor indicates that the router is down in the Dead Interval field. The dead interval is an unsigned integer. The default is four times the hello interval, or 40 seconds. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 to 8192.

Step 8 Enter the time in seconds between the hello packets that are sent on an interface in the Hello Interval field. The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 to 8192. The default is 10.

Step 9 Enter the time in seconds between LSA retransmissions for adjacencies that belong to the interface in the Retransmit Interval field. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 to 8192. The default is 5.

Step 10 Enter the estimated time in seconds that is required to send a link-state update packet on the interface in the Transmit Delay field. The integer value must be greater than zero. LSAs in the update packet have their own ages incremented by this amount before transmission. The range of values can be from 1 to 8192. The default is 1.

Step 11 In the Authentication area, check the **Enable Authentication** check box to enable authentication.

Step 12 Enter the security policy index, which must be a number from 256 to 4294967295, in the Security Policy Index field.

Step 13 Choose the authentication algorithm from the Authentication Algorithm drop-down list. Supported values are SHA-1 and MD5. When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.

Step 14 Enter the authentication key in the Authentication Key field. The key must include 32 hexadecimal characters.

Step 15 Choose the encryption algorithm from the Encryption Algorithm drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 16 Enter the encryption key in the Encryption Key field.

Step 17 Click **OK**.

Step 18 Click **Apply** to save your changes.

Configuring OSPFv3 Passive Interfaces

To configure OSPFv3 passive interfaces, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Passive Interfaces area allows you to enable passive OSPFv3 routing on an interface. Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface. In the Passive Interfaces area, choose the following settings:

- Check the **Global passive** check box to make all of the interfaces listed in the table passive. Uncheck individual interfaces to make them non-passive.
- Uncheck the **Global passive** check box to make all of the interfaces non-passive. Check individual interfaces to make them passive.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configuring OSPFv3 Administrative Distance

To configure OSPFv3 administrative distance for IPv6 routes, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Administrative Route Distances area allows you to modify the settings that were used to configure administrative route distances. The administrative route distance is an integer from 10 to 254. In the Administrative Route Distances area, enter the following values:

- The Inter Area, which specifies the inter-area routes for OSPV for IPv6 routes.
- The Intra Area, which specifies the intra-area routes for OSPF for IPv6 routes.
- The External, which specifies the external type 5 and type 7 routes for OSPF for IP6 routes.

Step 4 Click **OK**.

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

Configuring OSPFv3 Timers

You can set LSA arrival, LSA pacing, and throttling timers for OSPFv3.

To set the minimum interval at which the ASA accepts the same LSA from OSPFv3 neighbors, perform the following steps:

To configure LSA flood packet pacing, perform the following steps:

To change the interval at which OSPFv3 LSAs are collected into a group and refreshed, check summed, or aged, perform the following steps:

To configure LSA retransmission packet pacing, perform the following steps:

LSA and SPF throttling provide a dynamic mechanism to slow down LSA updates in OSPFv3 during times of network instability and allow faster OSPFv3 convergence by providing LSA rate limiting in milliseconds.

To configure LSA and SPF throttling timers, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA arrival, LSA pacing, LSA retransmission, LSA throttle, and SPF throttle times. In the Timers area, enter the following values:

- The LSA Arrival, which specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6000,000 milliseconds. The default is 1000 milliseconds.
- The LSA Flood Pacing, which specifies the time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds.
- The LSA Group Pacing, which specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- The LSA Retransmission Pacing, which specifies the time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds.
- The LSA Throttle Initial, which specifies the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 milliseconds.
- The LSA Throttle Min Hold, which specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.
- The LSA Throttle Max Wait, which specifies the maximum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.

Note For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- The SPF Throttle Initial, specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds.
- The SPF Throttle Min Hold, which specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds.
- The SPF Throttle Max Wait, which specifies the maximum wait time in milliseconds for SPF calculations. The default value is 10000 milliseconds.

Note For SPF throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Defining Static OSPFv3 Neighbors

You need to define static OSPFv3 neighbors to advertise OSPF routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv3 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Before you begin, you must create a static route to the OSPFv3 neighbor. See Chapter 31, "Configuring Static and Default Routes," for more information about creating static routes.

To define a static OSPFv3 neighbor, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add or Edit Static Neighbor dialog box appears. This 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, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPFv3 processes.
- You need to define a static route for each static neighbor.

Step 3 From the Interface drop-down list, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the Link-local Address field, enter the IPv6 address of the static neighbor.

Step 5 (Optional) In the Priority field, enter the priority level.

Step 6 (Optional) In the Poll Interval field, enter the poll interval in seconds.

Step 7 Click **OK**.

Sending Syslog Messages

To configure the router to send a syslog message when an OSPFv3 neighbor goes up or down, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Adjacency Changes area allows you to modify the settings for sending syslog messages when an OSPFv3 neighbor goes up or down. In the Adjacency Changes area, do the following:

- To send a syslog message when an OSPFv3 neighbor goes up or down, check the **Log Adjacency Changes** check box.
- To sends a syslog message for each state, not only when an OSPFv3 neighbor goes up or down, check the **Include Details** check box.

Step 4 Click **OK**.

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

Suppressing Syslog Messages

To suppress the sending of syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **Ignore LSA MOSPF** check box, then click **OK**.

Calculating Summary Route Costs

To calculate summary route costs according to RFC 1583, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **RFC1583 Compatible** check box, then click **OK**.

Generating a Default External Route into an OSPFv3 Routing Domain

To generate a default route into an OSPFv3 routing domain, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 In the Default Information Originate Area, do the following:

a. Check the **Enable** check box to enable the OSPFv3 routing process.

b. Check the **Always advertise** check box to always advertise the default route, whether or not one exists.

c. Enter the metric used for generating the default route in the Metric field. Valid metric values range from 0 to 16777214. The default value is 10.

d. From the Metric Type drop-down list, choose the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are the following:

- 1—Type 1 external route
- 2—Type 2 external route

The default is the Type 2 external route.

e. From the Route Map drop-down list, choose the routing process that generates the default route if the route map is satisfied.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configuring an IPv6 Summary Prefix

To configure an IPv6 summary prefix, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix**.

Step 2 To add a new summary prefix, click **Add**. To modify an existing summary prefix, click **Edit**. To remove a summary prefix, click **Delete**.

The Add Summary Prefix dialog box or Edit Summary Prefix dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 5 Check the **Advertise** check box to advertise routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.

Step 6 Enter the tag value that you can use as a match value for controlling redistribution through route maps in the Tag field.

Step 7 Click **OK**.

Step 8 Click **Apply** to save your changes.

Redistributing IPv6 Routes

To redistribute connected routes into an OSPFv3 process, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Redistribution**.

Step 2 To add new parameters for redistributing connected routes into an OSPFv3 process, click **Add**. To modify existing parameters for redistributing connected routes into an OSPFv3 process, click **Edit**. To remove a selected set of parameters, click **Delete**.

The Add Redistribution dialog box or Edit Redistribution dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the source protocol from which routes are being redistributed from the Source Protocol drop-down list. The supported protocols are connected, static, and OSPF.

Step 5 Enter the metric value in the Metric field. When redistributing routes from one OSPF process into another OSPF process on the same router, the metric is carried through from one process to the other if no metric value is specified. When redistributing other processes into an OSPF process, the default metric is 20 when no metric value is specified.

Step 6 Choose the metric type from the Metric Type drop-down list. The available options are None, 1, and 2.

Step 7 (Optional) Enter the tag value in the Tag field. This parameter specifies the 32-bit decimal value attached to each external route, which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.

Step 8 Choose the route map from the Route Map drop-down list to check for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported.

Step 9 To include connected routes in the redistribution, check the **Include connected** check box.

Step 10 Check the **Match** check box to redistribute routes into other routing domains, then check one of the following check boxes:

- **Internal** for routes that are internal to a specific autonomous system
- **External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 1 external routes
- **External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 2 external routes

- **NSSA External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 1 external routes
- **NSSA External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 2 external routes

Step 11 Click **OK**.

Step 12 Click **Apply** to save your changes.

Removing the OSPF Configuration

To remove the entire OSPFv2 configuration that you have already enabled, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Uncheck the **Enable this OSPF Process** check box.

Step 3 Click **Apply**.

To remove the entire OSPFv3 configuration that you have already enabled, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 Uncheck the **Enable OSPFv3 Process** check box.

Step 3 Click **Apply**.

Configuration Example for OSPFv2

The following example shows how to enable and configure OSPFv2 with various optional processes:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 2 Click the **Process Instances** tab and in the OSPF Process 1 field, type **2**.

Step 3 Click the **Area/Networks** tab, and click **Add**.

Step 4 Enter **0** in the Area ID field.

Step 5 In the Area Networks area, enter **10.0.0.0** in the IP Address field.

Step 6 Choose 255.0.0.0 from the Netmask drop-down list.

Step 7 Click **OK**.

Step 8 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Redistribution**.

Step 9 Click **Add**.

The Add/Edit OSPF Redistribution Entry dialog box appears.

Step 10 In the Protocol area, click the **OSPF** radio button to choose the source protocol from which the routes are being redistributed. Choosing OSPF redistributes routes from another OSPF routing process.

Step 11 Choose the OSPF process ID from the OSPF Process drop-down list.

Step 12 In the Match area, check the **Internal** check box.

Step 13 In the Metric Value field, enter **5** for the metric value for the routes being redistributed.

Step 14 From the Metric Type drop-down list, choose 1 for the Metric Type value.

Step 15 From the Route Map drop-down list, choose 1.

Step 16 Click **OK**.

Step 17 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Interface**.

Step 18 From the Properties tab, choose the **inside** interface and click **Edit.**

The Edit OSPF Properties dialog box appears.

Step 19 In the Cost field, enter **20**.

Step 20 Click **Advanced**.

Step 21 In the Retransmit Interval field, enter **15**.

Step 22 In the Transmit Delay field, enter **20**.

Step 23 In the Hello Interval field, enter **10**.

Step 24 In the Dead Interval field, enter **40**.

Step 25 Click **OK**.

Step 26 In the Edit OSPF Properties dialog box, enter **20** in the Priorities field, and click **OK**.

Step 27 Click the **Authentication** tab.

The Edit OSPF Authentication dialog box appears.

Step 28 In the Authentication area, click the **MD5** radio button.

Step 29 In the MD5 and Key ID area, enter **cisco** in the MD5 Key field, and **1** in the MD5 Key ID field.

Step 30 Click **OK**.

Step 31 Choose **Configuration > Device Setup > Routing > OSPF > Setup,** and click the **Area/Networks** tab.

Step 32 Choose the **OSPF 2** process and click **Edit**.

The Edit OSPF Area dialog box appears.

Step 33 In the Area Type area, choose **Stub**.

Step 34 In the Authentication area, choose **None**, and enter **20** in the Default Cost field.

Step 35 Click **OK**.

Step 36 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.

Step 37 Click the **Process Instances** tab and check the **OSPF process 2** check box.

Step 38 Click **Advanced**.

The Edit OSPF Area dialog box appears.

Step 39 In the Timers area, enter **10** in the SPF Delay Time field and **20** in the SPF Hold Time field.

Step 40 In the Adjacency Changes area, check the **Log Adjacency Change Details** check box.

Step 41 Click **OK**.

Step 42 Click **Reset**.

Configuration Example for OSPFv3

The following example shows how to configure OSPFv3 routing in ASDM:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

Step 2 On the Process Instances tab, do the following:

a. Check the **Enable OSPFv3 Process** check box.

b. Enter **1** in the Process ID field.

Step 3 Click the **Areas** tab. then click **Add** to display the Add OSPFv3 Area dialog box.

Step 4 From the OSPFv3 Process ID drop-down list, choose **1**.

Step 5 Enter **22** in the Area ID field.

Step 6 Choose **Normal** from the Area Type drop-down list.

Step 7 Enter **10** in the Default Cost field.

Step 8 Check the **Redistribution imports routes to normal and NSSA areas** check box.

Step 9 Enter **20** in the Metric field.

Step 10 Choose **1** from the Metric Type drop-down list.

Step 11 Check the **inside** check box as the specified interface being used.

Step 12 Check the **Enable Authentication** check box.

Step 13 Enter **300** in the Security Policy Index field.

Step 14 Choose **SHA-1** from the Authentication Algorithm drop-down list.

Step 15 Enter **12345ABCDE** in the Authentication Key field.

Step 16 Choose **DES** from the Encryption Algorithm drop-down list.

Step 17 Enter **1122334455aabbccddee** in the Encryption Key field.

Step 18 Click **OK**.

Step 19 Click the **Route Summarization** tab, then click **Add** to display the Add Route Summarization dialog box.

Step 20 Choose **1** from the Process ID drop-down list.

Step 21 Choose **22** from the Area ID drop-down list.

Step 22 Enter **2000:122::/64** in the IPv6 Prefix/Prefix Length field.

Step 23 (Optional) Enter **100** in the Cost field.

Step 24 Check the **Advertised** check box.

Step 25 Click **OK**.

Step 26 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Interface**.

Step 27 Click the **Properties** tab.

Step 28 Check the **inside** check box and click **Edit** to display the Edit OSPF Properties dialog box.

Step 29 In the Cost field, enter **20**.

Step 30 Enter **1** in the Priority field.

Step 31 Check the **point-to-point** check box.

Step 32 In the Dead Interval field, enter **40**.

Step 33 In the Hello Interval field, enter **10**.

Step 34 In the Retransmit Interval field, enter **15**.

Step 35 In the Transmit Delay field, enter **20**.

Step 36 Click **OK**.

Step 37 In the main ASDM window, choose **Configuration > Device Setup > Routing > Redistribution**.

Step 38 Choose **1** from the Process ID drop-down list.

Step 39 Choose **OSPF** from the Source Protocol drop-down list.

Step 40 Enter **50** in the Metric field.

Step 41 Choose **1** from the Metric Type drop-down list.

Step 42 Click **OK**.

Step 43 Click **Apply** to save your changes.

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can also use the information provided to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your device packets are taking through the network.

To monitor or display various OSPFv2 routing statistics in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Monitoring > Routing > OSPF LSAs**.

Step 2 You can select and monitor OSPF LSAs, Types 1 through 5 and 7. Each pane shows one LSA type, as follows:

- Type 1 LSAs represent the routes in an area under a process.
- Type 2 LSAs show the IP address of the designated router that advertises the routers.
- Type 3 LSAs show the IP address of the destination network.
- Type 4 LSAs show the IP address of the AS boundary router.
- Type 5 LSAs and Type 7 LSAs show the IP address of the AS external network.

Step 3 Click **Refresh** to update each LSA type pane.

Step 4 In the main ASDM window, choose **Monitoring > Routing > OSPF Neighbors**.

In the OSPF Neighbors pane, each row represents one OSPF neighbor. In addition, the OSPF Neighbors pane shows the network on which the neighbor is running, the priority, the state, the amount of dead time in seconds, the IP address of the neighbor, and the interface on which it is running. For a list of possible states for an OSPF neighbor, see RFC 2328.

Step 5 Click **Refresh** to update the OSPF Neighbors pane.

To monitor or display various OSPFv3 routing statistics in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Monitoring > Routing > OSPFv3 LSAs**.

Step 2 You can select and monitor OSPFv3 LSAs. Choose a link-state type to display its status according to specified parameters from the Link State type drop-down list. The supported link-state types are router, network, inter-area prefix, inter-area router, AS external, NSSA, link, and intra-area prefix.

Step 3 Click **Refresh** to update each link-state type.

Step 4 In the main ASDM window, choose **Monitoring > Routing > OSPFv3 Neighbors**.

In the OSPFv3 Neighbors pane, each row represents one OSPFv3 neighbor. In addition, the OSPFv3 Neighbors pane shows the IP address of the neighbor, the priority, the state, the amount of dead time in seconds, and the interface on which it is running. For a list of possible states for an OSPFv3 neighbor, see RFC 5340.

Step 5 Click **Refresh** to update the OSPFv3 Neighbors pane.

Additional References

For additional information related to implementing OSPF, see the following section:

- RFCs

RFCs

RFC	Title
2328	OSPFv2
4552	OSPFv3 Authentication
5340	OSPF for IPv6

Feature History for OSPF

Table 33-1 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 33-1 Feature History for OSPF

Feature Name	Platform Releases	Feature Information
OSPF Support	7.0(1)	Support was added for route data, authentication, and redistribution and monitoring of routing information using the Open Shortest Path First (OSPF) routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > OSPF.
Dynamic Routing in Multiple Context Mode	9.0(1)	OSPFv2 routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup
Clustering		For OSPFv2 and OSPFv3, bulk synchronization, route synchronization, and Layer 2 load balancing are supported in the clustering environment.
OSPFv3 Support for IPv6		OSPFv3 routing is supported for IPv6. We introduced the following screens: Configuration > Device Setup > Routing > OSPFv3 > Setup, Configuration > Device Setup > Routing > OSPFv3 > Interface, Configuration > Device Setup > Routing > OSPFv3 > Redistribution, Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix, Configuration > Device Setup > Routing > OSPFv3 > Virtual Link, Monitoring > Routing > OSPFv3 LSAs, Monitoring > Routing > OSPFv3 Neighbors.

Configuring RIP

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information, using the Routing Information Protocol (RIP).

This chapter includes the following sections:

- Information About RIP, page 34-1
- Licensing Requirements for RIP, page 34-3
- Guidelines and Limitations, page 34-3
- Configuring RIP, page 34-4
- Customizing RIP, page 34-5
- Monitoring RIP, page 34-12
- Configuration Example for RIP, page 34-12
- Feature History for RIP, page 34-13

Information About RIP

This section includes the following topics:

- Routing Update Process, page 34-2
- RIP Routing Metric, page 34-2
- RIP Stability Features, page 34-2
- RIP Timers, page 34-2
- Using Clustering, page 34-3

The Routing Information Protocol, or RIP, as it is more commonly called, is one of the most enduring of all routing protocols. RIP has four basic components: routing update process, RIP routing metrics, routing stability, and routing timers. Devices that support RIP send routing-update messages at regular intervals and when the network topology changes. These RIP packets include information about the networks that the devices can reach, as well as the number of routers or gateways that a packet must travel through to reach the destination address. RIP generates more traffic than OSPF, but is easier to configure.

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 ASA supports 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 ASA receives reliable routing information from a trusted source.

RIP has advantages over static routes because the initial configuration is simple, and you do not need to update the configuration when the topology changes. The disadvantage to RIP is that there is more network and processing overhead than in static routing.

Routing Update Process

RIP sends routing-update messages at regular intervals and when the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop. RIP routers maintain only the best route (the route with the lowest metric value) to a destination. After updating its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates are sent independently of the regularly scheduled updates that RIP routers send.

RIP Routing Metric

RIP uses a single routing metric (hop count) to measure the distance between the source and a destination network. Each hop in a path from source to destination is assigned a hop count value, which is typically 1. When a router receives a routing update that contains a new or changed destination network entry, the router adds 1 to the metric value indicated in the update and enters the network in the routing table. The IP address of the sender is used as the next hop.

RIP Stability Features

RIP prevents routing loops from continuing indefinitely by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of hops in a path is 15. If a router receives a routing update that contains a new or changed entry, and if increasing the metric value by 1 causes the metric to be infinity (that is, 16), the network destination is considered unreachable. The downside of this stability feature is that it limits the maximum diameter of a RIP network to less than 16 hops.

RIP includes a number of other stability features that are common to many routing protocols. These features are designed to provide stability despite potentially rapid changes in network topology. For example, RIP implements the split horizon and hold-down mechanisms to prevent incorrect routing information from being propagated.

RIP Timers

RIP uses numerous timers to regulate its performance. These include a routing-update timer, a route-timeout timer, and a route-flush timer. The routing-update timer clocks the interval between periodic routing updates. Generally, it is set to 30 seconds, with a small random amount of time added whenever the timer is reset. This is done to help prevent congestion, which could result from all routers

simultaneously attempting to update their neighbors. Each routing table entry has a route-timeout timer associated with it. When the route-timeout timer expires, the route is marked invalid but is retained in the table until the route-flush timer expires.

Using Clustering

For information about using clustering with RIP, see the "Dynamic Routing and Clustering" section on page 30-9.

Licensing Requirements for RIP

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 context mode only.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

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 ASA 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.

Limitations

RIP has the following limitations:

- The ASA 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 ASA.

Configuring RIP

This section describes how to enable and restart the RIP process on the ASA.

After you have enabled RIP, see the "Customizing RIP" section on page 34-5 to learn how to customize the RIP process on the ASA.

Note If you want to redistribute a route by defining which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process, you must first generate a default route. For information, see the "Configuring a Default Static Route" section on page 31-6 and then define a route map. For information, see the "Defining a Route Map" section on page 32-4.

Enabling RIP

You can only enable one RIP routing process on the ASA. After you enable the RIP routing process, you must define the interfaces that will participate in that routing process using the **network** command. By default, the ASA sends RIP Version 1 updates and accepts RIP Version 1 and Version 2 updates.

In ASDM, to enable a RIP process, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

The main RIP Setup pane appears.

From this pane, you can perform the following tasks:

- Enable Auto-summarization. See the "Configuring Route Summarization" section on page 34-8.
- Enable RIP version. See the "Configuring the RIP Version" section on page 34-5.
- Enable default information origination.
- Define an IP Address for a Network to Add. See the "Filtering Networks in RIP" section on page 34-8.
- Configure an Interface. See the "Configuring Interfaces for RIP" section on page 34-5.

Step 2 Check the **Enable RIP routing** check box.

After the Enable RIP routing box has been checked, you can enable RIP on the ASA and configure global RIP protocol parameters. You can only enable a single RIP process on the ASA. When you enable RIP, it is enabled on all interfaces. Checking this check box also enables the other fields in this pane. Uncheck this check box to disable RIP routing on the ASA.

Step 3 Click **Apply**.

Step 4 To customize the RIP process, see the "Configuring RIP" section on page 34-4.

Customizing RIP

This section describes how to configure RIP and includes the following topics:

- Configuring the RIP Version, page 34-5
- Configuring Interfaces for RIP, page 34-5
- Configuring the RIP Send and Receive Version on an Interface, page 34-7
- Configuring Route Summarization, page 34-8
- Filtering Networks in RIP, page 34-8
- Redistributing Routes into the RIP Routing Process, page 34-10
- Enabling RIP Authentication, page 34-11
- Restarting the RIP Process, page 34-12

Configuring the RIP Version

In ASDM, you can specify the version of RIP used by the ASA by performing the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box, and click **Apply**.

Step 3 Check the **Enable RIP version** check box.

Checking this check box specifies the version of RIP used by the ASA. If this check box is unchecked, then the ASA sends RIP Version 1 updates and accepts RIP Version 1 and Version 2 updates. This setting can be overridden on a per-interface basis in the Interface pane. For more information about configuring interfaces, see the "Configuring Interfaces for RIP" section on page 34-5. Indicate the version of RIP to be used by choosing one of the following:

- Version 1, which specifies that the ASA only sends and receives RIP Version 1 updates. Any Version 2 updates received are dropped.
- Version 2, which specifies that the ASA only sends and receives RIP Version 2 updates. Any Version 1 updates received are dropped.

Step 4 Click **Apply**.

Configuring Interfaces for RIP

If you have an interface that you do not want to have participate in RIP routing, but that is attached to a network that you want advertised, you can configure the network that includes the network to which the interface is attached, and configure the passive interfaces to prevent that interface from using RIP. Additionally, you can specify the version of RIP that is used by the ASA for updates.

In ASDM, you can configure an interface in RIP used by the ASA so that all interfaces on the ASA are set to passive RIP mode. The ASA 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, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box.

Step 3 In the Passive Interfaces area, check the check box in the Passive column for those interfaces that you want to have operate in passive mode. The other interfaces will still send and receive RIP broadcasts.

Step 4 Click **Apply**.

Note Individual interfaces can be made passive only if the global passive mode is not enabled. Uncheck the **Global Passive** check box to make individual interfaces passive using the Passive Interfaces table.

You can override this setting on a per-interface basis in the Interface pane. For more information, see the "Editing a RIP Interface" section on page 34-6.

Editing a RIP Interface

In ASDM, the Interface pane allows you to configure interface-specific RIP settings, such as the version of RIP that the interface sends and receives and the authentication method, if any, that are used for the RIP broadcasts.

To edit an interface that you have previously set up and configured, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box and click **Apply**.

Step 3 Choose **Configuration > Device Setup > Routing > RIP > Interfaces**.

Step 4 Click **Add** or **Edit**.

The Add or Edit RIP Interface Entry dialog box appears and allows you to configure the interface-specific RIP settings.

Step 5 (Optional) Choose the following options according to your preferences:

- **Override Global Send Version**—Check this check box to specify the RIP version sent by the interface. Choose one of 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. Choose one of 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 authentication. Specify the following settings:
  - Key, which is the key used by the authentication method, and can be up to 16 characters long.
  - Key ID, which is the key ID used by the authentication method. Valid values range from 0 to 255.
  - Authentication Mode—You can choose one of the following authentication modes:

    MD5 to use MD5 for RIP message authentication.

    Use clear text for RIP message authentication (not recommended).

Step 6 Click **Apply**.

Configuring the RIP Send and Receive Version on an Interface

You can override the globally-set version of RIP that the ASA uses to send and receive RIP updates on a per-interface basis.

To configure the RIP version for sending and receiving updates, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** box, and click **Apply**.

Step 3 Choose **Configuration > Device Setup > Routing > RIP > Interfaces**.

Step 4 Click **Edit**.

The Edit RIP Interface Entry dialog box appears, which allows you to configure the interface-specific RIP settings for sending and receiving.

Step 5 In the Send Version area, check the **Override global send version** check box to specify the RIP version sent by the interface. Choose one of the following:

- Version 1
- Version 2
- Version 1 & 2

  Unchecking this check box restores the global setting.

Step 6 In the Receive Version area, check the **Override global receive version** 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. Choose one of from the following:

- Version 1
- Version 2
- Version 1 & 2

  Unchecking this check box restores the global setting.

Step 7 Click **Apply**.

Configuring Route Summarization

Note RIP Version 1 always uses automatic route summarization. You cannot disable this feature for RIP Version 1. RIP Version 2 uses automatic route summarization by default.

The RIP routing process summarizes on network number boundaries, which can cause routing problems if you have noncontiguous networks.

For example, if you have a router with the networks 192.168.1.0, 192.168.2.0, and 192.168.3.0 connected to it, and those networks all participate in RIP, the RIP routing process creates the summary address 192.168.0.0 for those routes. If an additional router is added to the network with the networks 192.168.10.0 and 192.168.11.0, and those networks participate in RIP, they will also be summarized as 192.168.0.0. To prevent the possibility of traffic being routed to the wrong location, you should disable automatic route summarization on the routers that are creating conflicting summary addresses.

Because RIP Version 1 always uses automatic route summarization, and RIP Version 2 always uses automatic route summarization by default, when configuring automatic route summarization, you only need to disable it.

In ASDM, you can enable or disable automatic route summarization in a RIP process by performing the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box, and click **Apply**.

Step 3 Check the **Enable Auto-Summarization** check box.

Uncheck 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 route summarization for RIP Version 1. If you are using RIP Version 2, you can turn off automatic route summarization by unchecking this check box. Disable automatic route summarization if you must perform routing between disconnected subnets. When automatic route summarization is disabled, subnets are advertised.

Step 4 Click **Apply**.

Filtering Networks in RIP

To filter the networks received in updates, perform the following steps:

Note Before you begin, you must create a standard access list that permits the networks that you want the RIP process to allow in the routing table and denies the networks that you want the RIP process to discard.

In ASDM, you can configure filter rules that 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.

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box, and click **Apply**.

Step 3 Choose **Configuration > Device Setup > Routing > RIP > Filter Rules**.

Step 4 Click **Add** or **Edit**.

The Add or Edit Filter Rule dialog box appears, which allows you to create or edit filter rules that apply to all interfaces or to a specific interface.

Step 5 From the Direction drop-down list, choose the direction in which the filter should act.

Choosing In filters networks on incoming RIP updates. Additionally, only the Interface drop-down list is visible.

If you choose Out as the filter direction, skip to Step 8.

Step 6 Choose the Interface type from the Interface drop-down list.

This setting allows you to choose a specific interface for the filter rule, or you can choose the All Interfaces option to apply the filter to all interfaces.

Step 7 (Optional) Add a network rule by clicking **Add**. Skip to the "Adding or Editing a Filter Rule" section on page 34-9.

Step 8 Choose Out to filter networks from outgoing RIP updates. Additionally, the Interface and Routing Process drop-down list becomes visible.

- Click the **Interface** radio button to choose a specific interface for the filter rule from the Interface drop-down list, or click the **All Interfaces** option to apply the filter to all interfaces.
- Click the **Routing Process** radio button to activate the Routing process drop-down list. Choose from the following routing process types:
  - connected
  - static
  - OSPF
  - RIP
  - EIGRP

Step 9 (Optional) Add a network rule by clicking **Add**. Continue to the "Adding or Editing a Filter Rule" section on page 34-9.

Adding or Editing a Filter Rule

After you have configured a filter rule (see the "Filtering Networks in RIP" section on page 34-8), you can add or edit a network rule below the selected rule in the list by performing the following steps:

Step 1 After you have selected the direction or Interface type from Step 5 or Step 8 of the previous procedure, click **Add** or **Edit** in the Filtering Networks in RIP area.

The Network Rule dialog box appears.

Step 2 Choose the action from the Action drop-down list. The default is Permit.

- Choose Permit if the specified network is not filtered from incoming or outgoing RIP advertisements.
- Choose Deny if the specified network is to be filtered from incoming or outgoing RIP advertisements.

Step 3 Enter the IP address for the network being filtered, if different than what is displayed, in the IP Address field.

By default, the IP Address field displays the IP Address for the network being filtered.

Step 4 Enter the netmask, if different than what is displayed, in the Netmask field.

By default, the Netmask field displays the network mask applied to the IP address.

Step 5 Click **OK**.

Redistributing Routes into the RIP Routing Process

You can redistribute routes from the OSPF, EIGRP, static, and connected routing processes into the RIP routing process.

Note Before you begin this procedure, you must create a route map to further define which routes from the specified routing protocol are redistributed in to the RIP routing process. See Chapter 32, "Defining a Route Map," for more information about creating a route map.

In ASDM, you can display the routes that are being redistributed from other routing processes into the RIP routing process by performing the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Redistribution**.

The Redistribution pane displays the routes that are being redistributed from other routing processes into the RIP routing process.

Step 2 Click **Add** or **Edit**.

If you clicked **Add**, the Add Route Redistribution dialog box allows you to add a new redistribution rule. If you clicked **Edit**, the Edit Route Redistribution dialog box allows you to change an existing rule.

Step 3 In the Protocol area, choose the routing protocol to redistribute into the RIP routing process:

- Static, for static routes.
- Connected, for directly connected networks.
- OSPF and OSPF ID, for 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, for 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.

Step 4 In the Metrics area, check the **Configure Metric Type** check box to specify a metric for the redistributed routes. If not specified, the routes are assigned a default metric of 0. When the check box is checked, choose from one of the following available values:

- **Transparent** to cause the current route metric to be used.
- **Value** to assign a specific metric value. Valid values range from 0 to 16.

Step 5 In the Optional area, choose the route map from the Route Map drop-down list. This route map specifies the name of a route map that must be specified before the route can be redistributed into the RIP routing process. Click **Manage** to configure a specific route map. For more information about configuring route maps, see the "Adding or Editing a Route Map" section on page 32-4.

Step 6 In the Match area, choose specific types of OSPF routes to redistribute by checking the check box next to the route type. This area is not active unless OSPF has been chosen in the Protocol area.

If you do not check any route types, Internal, External 1, and External 2 routes are redistributed by default. The Match types are:

- Internal, in which routes internal to the AS are redistributed.
- External 1, in which Type 1 routes external to the AS are redistributed.
- External 2, in which Type 2 routes external to the AS are redistributed.
- NSSA External 1, in which Type 1 routes external to an NSSA are redistributed.
- NSSA External 2, in which Type 2 routes external to an NSSA are redistributed.

Step 7 Click **OK**.

Enabling RIP Authentication

Note The ASA supports RIP message authentication for RIP Version 2 messages.

RIP route authentication provides MD5 authentication of routing updates from the RIP routing protocol. The MD5 keyed digest in each RIP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

RIP route authentication is configured on a per-interface basis. All RIP neighbors on interfaces configured for RIP message authentication must be configured with the same authentication mode and key for adjacencies to be established.

Note Before you can enable RIP route authentication, you must enable RIP.

To enable RIP authentication on an interface, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box, and click **Apply**. If you uncheck this check box, the ASA sends RIP Version 1 updates and accepts RIP Version 1 and Version 2 updates. You can override this setting on a per-interface basis in the Interface pane. Version 1 specifies that the ASA only sends and receives RIP Version 1 updates. Any Version 2 updates received are dropped. Version 2 specifies that the ASA only sends and receives RIP Version 2 updates. Any Version 1 updates received are dropped.

Step 3 Choose **Configuration > Device Setup > Routing > RIP > Interface**.

Step 4 Click **Edit**.

The Edit RIP Interface Entry dialog box appears, which allows you to configure the interface-specific RIP settings.

Step 5 In the Authentication area, check the **Enable Authentication** check box to enable RIP authentication. Uncheck this check box to disable RIP authentication.

Step 6 In the Key field, enter the key used by the authentication method. This entry can include up to 16 characters.

Step 7 In the Key ID field, enter the key ID. Valid values range from 0 to 255.

Step 8 Choose the type of authentication mode that you want to use by clicking the radio button next to one of the following:

- **MD5** to use MD5 for RIP message authentication.
- **cleartext** to use cleartext for RIP message authentication (not recommended).

Step 9 Click **Apply**.

Restarting the RIP Process

To remove the entire RIP configuration, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Click **Reset**.

Monitoring RIP

To monitor or display various RIP routing statistics in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Monitoring > Routing > Routes**.

Step 2 From this pane, you can choose to monitor the following:

- **IPv4**
- **IPv6**
- **Both**

Configuration Example for RIP

The following example shows how to enable and configure RIP with various optional processes:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Check the **Enable RIP routing** check box and click **Apply**.

Step 3 Check the **Enable default information originate** check box.

For more information about defining a route map, see the "Defining a Route Map" section on page 32-4.

Step 4 Check the **Enable RIP version** check box and choose **Version 1**.

Step 5 In the Networks area, enter **225.25.24.225** in the IP Network to Add field.

Step 6 In the Passive Interface area, click the check box next to the interface that you want to be passive in the Passive Interfaces table.

Step 7 Click **Apply**.

Step 8 Choose **Configuration > Device Setup > Routing > RIP > Redistribution**.

Step 9 Click **Edit**.

Step 10 In the Protocol area, choose **Connected**.

Step 11 In the Metric area, check the **Configure Metric Type** check box and choose **Transparent Mode** (default).

Step 12 In the Optional area, choose a route map from the Route Map drop-down list.

Step 13 Click **Manage** to configure a specific route map. For more information about configuring route maps, see the "Adding or Editing a Route Map" section on page 32-4.

Step 14 Click **OK**.

Feature History for RIP

Table 34-1 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 34-1 Feature History for RIP

Feature Name	Releases	Feature Information
RIP support	7.0(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Routing Information Protocol (RIP). We introduced the following screen: Configuration > Device Setup > Routing > RIP.
Clustering	9.0(1)	For RIP, bulk synchronization, route synchronization, and layer 2 load balancing are supported in the clustering environment.

Configuring Multicast Routing

This chapter describes how to configure the ASA to use the multicast routing protocol and includes the following sections:

- Information About Multicast Routing, page 35-1
- Licensing Requirements for Multicast Routing, page 35-3
- Guidelines and Limitations, page 35-3
- Enabling Multicast Routing, page 35-3
- Customizing Multicast Routing, page 35-4
- Configuration Example for Multicast Routing, page 35-17
- Additional References, page 35-18
- Feature History for Multicast Routing, page 35-19

Information About Multicast Routing

Multicast routing is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast routing include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Multicast routing protocols delivers source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. Multicast packets are replicated in the network by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols resulting in the most efficient delivery of data to multiple receivers possible.

The ASA supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single ASA.

Note The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

This section includes the following topics:

- Stub Multicast Routing, page 35-2
- PIM Multicast Routing, page 35-2

- Multicast Group Concept, page 35-2
- Clustering, page 35-2

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the ASA acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the ASA forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the ASA cannot be configured for PIM.

The ASA supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees rooted at a single Rendezvous Point per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bi-directional PIM is a variant of PIM-SM that builds bi-directional shared trees connecting multicast sources and receivers. Bi-directional trees are built using a DF election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point, and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during Rendezvous Point discovery and provides a default route to the Rendezvous Point.

Note If the ASA is the PIM RP, use the untranslated outside address of the ASA as the RP address.

Multicast Group Concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries—the hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using IGMP. Hosts must be a member of the group to receive the data stream. For information about how to configure multicast groups, see the "Configuring a Multicast Group" section on page 35-14.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Layer 2 clustering, the master unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, slave units may forward multicast data packets. All data flows are full flows. Stub

forwarding flows are also supported. Because only one unit receives multicast packets in Layer 2 clustering, redirection to the master unit is common. In Layer 3 clustering, units do not act independently. All data and routing packets are processed and forwarded by the master unit. Slave units drop all packets that have been sent.

For more information about clustering, see Chapter 10, "Configuring a Cluster of ASAs."

Licensing Requirements for Multicast Routing

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 context mode. In multiple context mode, unshared interfaces and shared interfaces are not supported.

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

In clustering, for IGMP and PIM, this feature is only supported on the master unit.

Enabling Multicast Routing

Enabling multicast routing lets you enable multicast routing on the ASA. 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.

To enable multicast routing, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast** routing check box.

Checking this check box enables IP multicast routing on the ASA. Unchecking this check box disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces. You can disable multicast on a per-interface basis.

Table 35-1 lists the maximum number of entries for specific multicast tables based on the amount of RAM on the ASA. Once these limits are reached, any new entries are discarded.

Table 35-1 Entry Limits for Multicast Tables

Table	16 MB	128 MB	128+ MB
MFIB	1000	3000	5000
IGMP Groups	1000	3000	5000
PIM Routes	3000	7000	12000

Customizing Multicast Routing

This section describes how to customize multicast routing and includes the following topics:

- Configuring Stub Multicast Routing and Forwarding IGMP Messages, page 35-4
- Configuring a Static Multicast Route, page 35-5
- Configuring IGMP Features, page 35-6
- Configuring PIM Features, page 35-10
- Configuring a Multicast Group, page 35-14
- Configuring a Bidirectional Neighbor Filter, page 35-15
- Configuring a Multicast Boundary, page 35-16

Configuring Stub Multicast Routing and Forwarding IGMP Messages

Note Stub multicast routing and PIM are not supported concurrently.

An ASA acting as the gateway to the stub area does not need to participate in PIM. Instead, you can configure it to act as an IGMP proxy agent and forward IGMP messages from hosts connected on one interface to an upstream multicast router on another interface. To configure the ASA as an IGMP proxy agent, forward the host join and leave messages from the stub area interface to an upstream interface.

To forward the host join and leave messages, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast routing** check box.

Step 3 Click **Apply** to save your changes.

Step 4 Choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

Step 5 To modify the specific interface from which you want to forward IGMP messages, select the interface and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 6 From the **Forward Interface** drop-down list, choose the specific interface from which you want to forward IGMP messages.

Step 7 Click **OK** to close this dialog box, then click **Apply** to save your changes.

Configuring a Static Multicast Route

Configuring 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.

When using PIM, the ASA expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

To configure a static multicast route or a static multicast route for a stub area, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > MRoute**.

Step 2 Choose **Add** or **Edit**.

The Add or Edit Multicast Route dialog box appears.

Use the Add Multicast Route dialog box to add a new static multicast route to the ASA. Use the Edit Multicast Route dialog box to change an existing static multicast route.

Step 3 In the Source Address field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.

Step 4 Choose the network mask for the IP address of the multicast source from the Source Mask drop-down list.

Step 5 In the Incoming Interface area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:

- In the Source Interface field, choose the incoming interface for the multicast route from the drop-down list.
- In the Destination Interface field, choose the destination interface that the route is forwarded through from the drop-down list.

Note You can specify the interface or the RPF neighbor, but not both at the same time.

Step 6 In the Administrative Distance field, choose 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.

Step 7 Click **OK**.

Configuring IGMP Features

IP hosts use the Internet Group Management Protocol (IGMP) to report their group memberships to directly connected multicast routers.

IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range of 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.

When you enable multicast routing on the ASA, IGMP Version 2 is automatically enabled on all interfaces.

Note Only the **no igmp** command appears in the interface configuration when you use the **show run** command. If the **multicast-routing** command appears in the device configuration, then IGMP is automatically enabled on all interfaces.

This section describes how to configure optional IGMP setting on a per-interface basis and includes the following topics:

- Disabling IGMP on an Interface, page 35-6
- Configuring IGMP Group Membership, page 35-7
- Configuring a Statically Joined IGMP Group, page 35-7
- Controlling Access to Multicast Groups, page 35-8
- Limiting the Number of IGMP States on an Interface, page 35-9
- Modifying the Query Messages to Multicast Groups, page 35-9
- Changing the IGMP Version, page 35-10

Disabling IGMP on an Interface

You can disable IGMP on specific interfaces. This information is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the ASA from sending host query messages on that interface.

To disable IGMP on an interface, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

The Protocol pane displays the IGMP parameters for each interface on the ASA.

Step 2 Choose the interface that you want to disable and click **Edit**.

Step 3 To disable the specified interface, uncheck the **Enable IGMP** check box.

Step 4 Click **OK**.

The Protocol pane displays Yes if IGMP is enabled on the interface, or No if IGMP is disabled on the interface.

Configuring IGMP Group Membership

You can configure the ASA to be a member of a multicast group. Configuring the ASA to join a multicast group causes upstream routers to maintain multicast routing table information for that group and keep the paths for that group active.

Note If you want to forward multicast packets for a specific group to an interface without the ASA accepting those packets as part of the group, see the "Configuring a Statically Joined IGMP Group" section on page 35-7.

To have the ASA join a multicast group, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Join Group**.

The Join Group pane appears.

Step 2 Click **Add** or **Edit**.

The Add IGMP Join Group dialog box allows you to configure an interface to be a member of a multicast group. The Edit IGMP Join Group dialog box allows you to change existing membership information.

Step 3 In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

Step 4 In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 5 Click **OK**.

Configuring a Statically Joined IGMP Group

Sometimes a group member cannot report its membership in the group because of some configuration, or there may be no members of a group on the network segment. However, you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group.

In the main ASDM window, choose **Configuration > Routing > Multicast > IGMP > Static Group** to configure the ASA to be a statically connected member of a group. With this method, the ASA 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 this interface is not a member of the multicast group.

To configure a statically joined multicast group on an interface, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Static Group**.

The Static Group pane appears.

Step 2 Click **Add** or **Edit**.

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.

Step 3 In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

Step 4 In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 5 Click **OK**.

Controlling Access to Multicast Groups

To control the multicast groups that hosts on the ASA interface can join, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Access Group**.

The Access Group pane appears. The table entries in the Access Group pane 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 or Edit Access Group dialog box for the selected entry.

Step 2 Click **Add** or **Edit**.

The Add Access Group or Edit Access Group dialog box appears. 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 dimmed when editing existing entries.

Step 3 Choose the interface name with which the access group is associated from the Interface drop-down list. You cannot change the associated interface when you are editing an existing access group.

Step 4 Choose permit from the Action drop-down list to allow the multicast group on the selected interface. Choose deny from the Action drop-down list to filter the multicast group from the selected interface.

Step 5 In the Multicast Group Address field, enter the address of the multicast group to which the access group applies.

Step 6 Enter the network mask for the multicast group address, or choose one of the common network masks from the Netmask drop-down list.

Step 7 Click **OK**.

Limiting the Number of IGMP States on an Interface

You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports exceeding the configured limits are not entered in the IGMP cache, and traffic for the excess membership reports is not forwarded.

To limit the number of IGMP states on an interface, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

Step 2 Choose the interface you want to limit from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 3 In the Group Limit field, enter the maximum number of host that can join on an interface. Valid values range from 0 to 500. The default value is 500. Setting this value to 0 prevents learned groups from being added, but manually defined memberships are still permitted.

Step 4 Click **OK**.

Modifying the Query Messages to Multicast Groups

The ASA sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the ASA. If the ASA discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packet for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the ASA does not receive a response to a host query within this amount of time, it deletes the group.

To change the query interval, query response time, and query timeout value, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

Step 2 Choose the interface you want to limit from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 3 In the Query Interval field, 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.

```
If the ASA does not hear a query message on an interface for the specified timeout value,
then the ASA becomes the designated router and starts sending the query messages.
```

Step 4 In the Query Timeout field, enter the period of time, in seconds, before which the ASA takes over as the requester for the interface after the previous requester has stopped doing so. Valid values range from 60 to 300 seconds. The default value is 255 seconds.

Step 5 Click **OK**.

Changing the IGMP Version

By default, the ASA runs IGMP Version 2, which enables several additional features.

All multicast routers on a subnet must support the same version of IGMP. The ASA does not automatically detect Version 1 routers and switch to Version 1. However, a mix of IGMP Version 1 and 2 hosts on the subnet works; the ASA running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

To control which version of IGMP is running on an interface,perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

Step 2 Choose the interface whose version of IGMP you want to change from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Interface dialog box appears.

Step 3 Choose the version number from the Version drop-down list.

Step 4 Click **OK**.

Configuring PIM Features

Routers use PIM to maintain forwarding tables for forwarding multicast diagrams. When you enable multicast routing on the ASA, PIM and IGMP are automatically enabled on all interfaces.

Note PIM is not supported with PAT. The PIM protocol does not use ports, and PAT only works with protocols that use ports.

This section describes how to configure optional PIM settings and includes the following topics:

- Enabling and Disabling PIM on an Interface, page 35-10
- Configuring a Static Rendezvous Point Address, page 35-11
- Configuring the Designated Router Priority, page 35-12
- Configuring and Filtering PIM Register Messages, page 35-12
- Configuring PIM Message Intervals, page 35-13
- Configuring a Route Tree, page 35-13
- Filtering PIM Neighbors, page 35-14

Enabling and Disabling PIM on an Interface

You can enable or disable PIM on specific interfaces. To enable or disable PIM on an interface, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Protocol**.

Step 2 Choose the interface on which you want to enable PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 Check the **Enable PIM** check box. To disable PIM, uncheck this check box.

Step 4 Click **OK**.

Configuring a Static Rendezvous Point Address

All routers within a common PIM sparse mode or bidir domain require knowledge of the PIM RP address. The address is statically configured using the **pim rp-address** command.

Note The ASA does not support Auto-RP or PIM BSR

You can configure the ASA to serve as RP to more than one group. The group range specified in the access list determines the PIM RP group mapping. If an access list is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4).

To configure the address of the PIM PR, perform the following steps:

Note The ASA always advertises the bidirectional capability in the PIM hello messages, regardless of the actual bidirectional configuration.

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Rendezvous Points**.

Step 2 Click **Add** or **Edit**.

The Add or Edit Rendezvous Point dialog box appears. 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. Additionally, you can click **Delete** to remove the selected multicast group entry from the table.

These restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Step 3 In the Rendezvous Point Address field, enter the IP address for the RP.

When editing an existing RP entry, you cannot change this value.

Step 4 Check the **Use bi-directional forwarding** check box if the specified multicast groups are to operate in bidirectional mode. The Rendezvous Point pane displays Yes if the specified multicast groups are to operate in bidirectional mode and displays No if the specified groups are to operate in sparse mode. In bidirectional mode, if the ASA receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.

Step 5 Click the **Use this RP for All Multicast Groups** radio button to use the specified RP for all multicast groups on the interface, or the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.

For more information about multicast groups, see the "Configuring a Multicast Group" section on page 35-14.

Step 6 Click **OK**.

Configuring the Designated Router Priority

The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, selecting the DR is based on the DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR.

By default, the ASA has a DR priority of 1. To change this value, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Protocol**.

Step 2 Choose the interface that you want to enable for PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 In the DR Priority field, type the value for the designated router priority for the selected interface. The router with the highest DR priority on the 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 ASA interface ineligible to become the default router.

Step 4 Click **OK**.

Configuring and Filtering PIM Register Messages

When the ASA is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. The Request Filter pane lets you define the multicast sources from which the ASA will accept PIM register messages.

To filter PIM register messages, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Request Filter**.

Step 2 Click **Add**.

The Request Filter Entry dialog box lets you define the multicast sources that are allowed to register with the ASA when the ASA acts as an RP. You create the filter rules based on the source IP address and the destination multicast address.

Step 3 From the Action drop-down list, choose Permit to create a rule that allows the specified source of the specified multicast traffic to register with the ASA, or choose Deny to create a rule that prevents the specified source of the specified multicast traffic from registering with the ASA.

Step 4 In the Source IP Address field, type the IP address for the source of the register message.

Step 5 In the Source Netmask field, type or choose the network mask from the drop-down list for the source of the register message.

Step 6 In the Destination IP Address field, type the multicast destination address.

Step 7 In the Destination Netmask field, type or choose the network mask from the drop-down list for the multicast destination address.

Step 8 Click **OK**.

Configuring PIM Message Intervals

Router query messages are used to select the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. Additionally, every 60 seconds, the ASA sends PIM join or prune messages.

To change these intervals, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Protocol**.

Step 2 Choose the interface that you want to enable for PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 In the Hello Interval field, type the frequency, in seconds, at which the interface sends PIM hello messages.

Step 4 In the Prune Interval field, type the frequency, in seconds, at which the interface sends PIM join and prune advertisements.

Step 5 Click **OK**.

Configuring a Route Tree

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the ASA should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

To configure a PIM leaf router tree, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Route Tree**.

Step 2 Click one of the following radio buttons:

- **Use Shortest Path Tree for All Groups**—Choose this option to use the shortest-path tree for all multicast groups.
- **Use Shared Tree for All Groups**—Choose this option to use the shared tree for all multicast groups.
- **Use Shared Tree for the Groups specified below**—Choose this option to use the shared tree for the groups specified in the Multicast Groups table. The shortest-path tree is used for any group that is not specified in the Multicast Groups table.

The Multicast Groups table displays the multicast groups to use with the shared tree.

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.

To edit a multicast group, see the "Configuring a Multicast Group" section on page 35-14.

Configuring a Multicast Group

Multicast groups are lists of access rules that define which multicast addresses are part of a group. A multicast group can include 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.

To configure a multicast group, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Rendezvous Points**.

Step 2 The Rendezvous Point pane appears. Click the group that you want to configure.

The Edit Rendezvous Point dialog box appears.

Step 3 Click the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.

Step 4 Click **Add** or **Edit**.

The Add or Edit Multicast Group dialog box appears.

Step 5 From the Action drop-down list, choose Permit to create a group rule that allows the specified multicast addresses, or choose Deny to create a group rule that filters the specified multicast addresses.

Step 6 In the Multicast Group Address field, type the multicast address associated with the group.

Step 7 From the Netmask drop-down list, choose the network mask for the multicast group address.

Step 8 Click **OK**.

Filtering PIM Neighbors

You can define the routers that can become PIM neighbors. By filtering the routers that can become PIM neighbors, you can do the following:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

To define neighbors that can become a PIM neighbor, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Neighbor Filter**.

Step 2 Choose the PIM neighbor that you want to configure from the table by clicking **Add/Edit/Insert**.

The Add/Edit/Insert Neighbor Filter Entry dialog box appears. The Add/Edit/Insert Neighbor Filter Entry dialog box lets you create the ACL entries for the multicast boundary ACL. You can also delete a selected PIM neighbor entry.

Step 3 Choose the interface name from the Interface Name drop-down list.

Step 4 From the Action drop-down list, choose Permit or Deny for the neighbor filter ACL entry.

Choosing Permit allows the multicast group advertisements to pass through the interface. Choosing 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.

Step 5 In the IP Address text field, enter the IP address of the multicast PIM group being permitted or denied. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 6 From the Netmask drop-down list, choose the netmask for the multicast group address.

Step 7 Click **OK**.

Configuring a Bidirectional Neighbor Filter

The Bidirectional Neighbor Filter pane shows the PIM bidirectional neighbor filters, if any, that are configured on the ASA. A PIM bidirectional neighbor filter 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 the DF election process.

When a PIM bidirectional neighbor filter configuration is applied to the ASA, an ACL appears in the running configuration with the name *interface-name*_multicast, in which the *interface-name* is the name of the interface to which the multicast boundary filter is applied. 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 ASA.

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 the 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 bidirectionally capable. Therefore, the following is true:

- If a permitted neighbor does not support bidir, then the DF election does not occur.
- If a denied neighbor supports bidir, then the DF election does not occur.
- If a denied neighbor does not support bidir, the DF election can occur.

To define the neighbors that can become a PIM bidirectional neighbor filter, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Bidirectional Neighbor Filter**.

Step 2 Double-click an entry in the PIM Bidirectional Neighbor Filter table to open the Edit Bidirectional Neighbor Filter Entry dialog box for that entry.

Step 3 Choose the PIM neighbor that you want to configure from the table by clicking **Add/Edit/Insert**.

The Add/Edit/Insert Bidirectional Neighbor Filter Entry dialog box appears, which lets you create ACL entries for the PIM bidirectional neighbor filter ACL

Step 4 Choose the interface name from the Interface Name drop-down list. Select the interface for which you are configuring the PIM bidirectional neighbor filter ACL entry.

Step 5 From the Action drop-down list, choose Permit or Deny for the neighbor filter ACL entry.

Choose Permit to allow the specified devices to participate in the DF election process. Choose Deny to prevent the specified devices from participating in the DF election process.

Step 6 In the IP Address text field, enter the IP address of the multicast PIM group being permitted or denied. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 7 From the Netmask drop-down list, choose the netmask for the multicast group address.

Step 8 Click **OK**.

Configuring a Multicast Boundary

Address scoping defines domain boundaries so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary on an interface for multicast group addresses by choosing **Configuration > Routing > Multicast > MBoundary** in ASDM. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

A standard ACL defines the range of affected addresses. When a boundary is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

To configure a multicast boundary, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Routing > Multicast > MBoundary**.

The MBoundary pane lets you configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

Step 2 Click **Edit**.

The Edit Boundary Filter dialog box appears and 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 ASA, the ACL appears in the running configuration with the name *interface-name*_multicast, where the *interface-name* is the name of the interface to which the multicast boundary filter is applied. If an ACL with that name already exists, a number is appended to the name (for example, inside_multicast_1).

Step 3 Choose the interface for which you are configuring the multicast boundary filter ACL from the Interface drop-down list.

Step 4 Check the **Remove any Auto-RP group range** check box to filter Auto-RP messages from sources denied by the boundary ACL. If the **Remove any Auto-RP group range** check box is unchecked, all Auto-RP messages are passed.

Step 5 Click **OK**.

Configuration Example for Multicast Routing

The following example shows how to enable and configure multicast routing with various optional processes:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast** routing check box, and click **Apply**.

Step 3 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > MRoute**.

Step 4 Click **Add** or **Edit**.

The Add or Edit Multicast Route dialog box appears.

Use the Add Multicast Route dialog box to add a new static multicast route to the ASA. Use the Edit Multicast Route dialog box to change an existing static multicast route.

Step 5 In the Source Address field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.

Step 6 Choose the network mask for the IP address of the multicast source from the Source Mask drop-down list.

Step 7 In the Incoming Interface area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:

- In the Source Interface field, choose the incoming interface for the multicast route from the drop-down list.
- In the Destination Interface field, choose the destination interface to which the route is forwarded through the selected interface from the drop-down list.

Note You can specify the interface or the RPF neighbor, but not both at the same time.

Step 8 In the Administrative Distance field, choose 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.

Step 9 Click **OK**.

Step 10 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Join Group**.

The Join Group pane appears.

Step 11 Click **Add** or **Edit**.

The Add IGMP Join Group dialog box allows you to configure an interface to be a member of a multicast group. The Edit IGMP Join Group dialog box allows you to change existing membership information.

Step 12 In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

Step 13 In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 14 Click **OK**.

Additional References

For additional information related to routing, see the following sections:

- Related Documents, page 35-19
- RFCs, page 35-19

Related Documents

Related Topic	Document Title
Technical details about the IGMP and multicast routing standards used for implementing the SMR feature	IETF draft-ietf-idmr-igmp-proxy-01.txt

RFCs

RFC	Title
RFC 2113	IP Router Alert Option
RFC 2236	IGMPv2
RFC 2362	PIM-SM
RFC 2588	IP Multicast and Firewalls

Feature History for Multicast Routing

Table 35-2 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 35-2 Feature History for Multicast Routing

Feature Name	Platform Releases	Feature Information
Multicast routing support	7.0(1)	Support was added for multicast routing data, authentication, and redistribution and monitoring of routing information using the multicast routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > Multicast.
Clustering support	9.0(1)	Support was added for clustering.

Configuring EIGRP

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP).

This chapter includes the following sections:

- Information About EIGRP, page 36-1
- Licensing Requirements for EIGRP, page 36-2
- Guidelines and Limitations, page 36-3
- Task List to Configure an EIGRP Process, page 36-3
- Configuring EIGRP, page 36-4
- Customizing EIGRP, page 36-6
- Monitoring EIGRP, page 36-18
- Feature History for EIGRP, page 36-19

Information About 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. Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet mask, support for partial updates, and support for multiple network layer protocols.

A router running EIGRP stores all the neighbor routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found. Its support for variable-length subnet masks permits routes to be automatically summarized on a network number boundary. In addition, EIGRP can be configured to summarize on any bit boundary at any interface. EIGRP does not make periodic updates. Instead, it sends partial updates only when the metric for a route changes. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

Neighbor discovery is the process that the ASA uses to dynamically learn of other routers on directly attached networks. EIGRP routers send out multicast hello packets to announce their presence on the network. When the ASA receives a hello packet from a new neighbor, it sends its topology table to the neighbor with an initialization bit set. When the neighbor receives the topology update with the initialization bit set, the neighbor sends its topology table back to the ASA.

The hello packets are sent out as multicast messages. No response is expected to a hello message. The exception to this is for statically defined neighbors. If you use the **neighbor** command, or configure the Hello Interval in ASDM, to configure a neighbor, the hello messages sent to that neighbor are sent as unicast messages. Routing updates and acknowledgements are sent out as unicast messages.

Once this neighbor relationship is established, routing updates are not exchanged unless there is a change in the network topology. The neighbor relationship is maintained through the hello packets. Each hello packet received from a neighbor includes a hold time. This is the time in which the ASA can expect to receive a hello packet from that neighbor. If the ASA does not receive a hello packet from that neighbor within the hold time advertised by that neighbor, the ASA considers that neighbor to be unavailable.

The EIGRP protocol uses four key algorithm technologies, four key technologies, including neighbor discovery/recovery, Reliable Transport Protocol (RTP), and DUAL, which is important for route computations. DUAL saves all routes to a destination in the topology table, not just the least-cost route. The least-cost route is inserted into the routing table. The other routes remain in the topology table. If the main route fails, another route is chosen from the feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination. The feasibility calculation guarantees that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation must occur. During route recomputation, DUAL queries the EIGRP neighbors for a route, who in turn query their neighbors. Routers that do no have a feasible successor for the route return an unreachable message.

During route recomputation, DUAL marks the route as active. By default, the ASA waits for three minutes to receive a response from its neighbors. If the ASA does not receive a response from a neighbor, the route is marked as stuck-in-active. All routes in the topology table that point to the unresponsive neighbor as a feasibility successor are removed.

Note EIGRP neighbor relationships are not supported through the IPsec tunnel without a GRE tunnel.

Using Clustering

For information about using clustering with EIGRP, see the "Dynamic Routing and Clustering" section on page 30-9.

Licensing Requirements for EIGRP

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 only in routed firewall mode. Transparent firewall mode is not supported.

Failover Guidelines

Supports Stateful Failover in single and multiple context mode.

IPv6 Guidelines

Does not support IPv6.

Clustering Guidelines

- Supports Layer 2 and Layer 3 clustering when configured to use both EIGRP and OSPFv2.
- In a Layer 3 cluster setup, EIGRP adjacencies can only be established between two contexts on a shared interface on the master unit. You can manually configure multiple neighbor statements corresponding to each cluster node separately to work around this issue.

Additional Guidelines

EIGRP instances cannot form adjacencies with each other across shared interfaces because inter-context exchange of multicast traffic is not supported.

Task List to Configure an EIGRP Process

To configure EIGRP routing on the ASA, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP**.

Step 2 Enable the EIGRP routing process by checking the **Enable this EIGRP process** check box on the Process Instances tab. See the "Enabling EIGRP" section on page 36-4 or the "Enabling EIGRP Stub Routing" section on page 36-5.

Step 3 Define the networks and interfaces that will participate in EIGRP routing on the Setup > Networks tab. See the "Defining a Network for an EIGRP Routing Process" section on page 36-7 for more information.

Step 4 (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 the "Filtering Networks in EIGRP" section on page 36-13 for more information.

Step 5 (Optional) Define route redistribution in the Redistribution pane.

You can redistribute routes discovered by RIP and OSPF to the EIGRP routing process. You can also redistribute static and connected routes to the EIGRP routing process. See the "Redistributing Routes Into EIGRP" section on page 36-12 for more information.

Step 6 (Optional) Define static EIGRP neighbors on the Static Neighbor pane.

See the "Defining an EIGRP Neighbor" section on page 36-11 for more information.

Step 7 (Optional) Define summary addresses on the Summary Address pane.

See the "Configuring the Summary Aggregate Addresses on Interfaces" section on page 36-9 for more information about defining summary addresses.

Step 8 (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 the "Configuring Interfaces for EIGRP" section on page 36-7 for more information.

Step 9 (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 the "Configuring Default Information in EIGRP" section on page 36-16 for more information.

Configuring EIGRP

This section describes how to enable the EIGRP process on your system. After you have enabled EIGRP, see the following sections to learn how to customize the EIGRP process on your system.

- Enabling EIGRP, page 36-4
- Enabling EIGRP Stub Routing, page 36-5

Enabling EIGRP

You can only enable one EIGRP routing process on the ASA.

To enable EIGRP,perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

The three tabs on the main EIGRP Setup pane used to enable EIGRP are as follows:

- The Process Instances tab lets you enable an EIGRP routing process for each context. Single context mode and multiple context mode are both supported. See the "Enabling EIGRP" section on page 36-4 and the "Enabling EIGRP Stub Routing" section on page 36-5 for more information.
- The Networks 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. See the "Defining a Network for an EIGRP Routing Process" section on page 36-7 for more information.
- The Passive Interfaces 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 that is configured as a passive interface.

Step 2 Check the **Enable this EIGRP process** check box.

You can only enable one EIGRP routing process on the device. You must enter an autonomous system number (AS) for the routing process in the EIGRP Process field before you can save your changes.

Step 3 In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 (Optional) Click **Advanced** to configure the EIGRP process settings, such as the router ID, default metrics, stub routing, neighbor changes, and the administrative distances for the EIGRP routes.

Step 5 Click the **Networks** tab.

Step 6 To add a new network entry, click **Add**.

The Add EIGRP Network dialog box appears. To remove a network entry, choose an entry in the table and click **Delete**.

Step 7 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 8 Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.

Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 9 Enter a network mask to apply to the IP address in the Network Mask field.

Step 10 Click **OK**.

Enabling EIGRP Stub Routing

You can enable, and configure the ASA as an EIGRP stub router. Stub routing decreases memory and processing requirements on the ASA. As a stub router, the ASA does not need to maintain a complete EIGRP routing table because it forwards all nonlocal traffic to a distribution router. Generally, the distribution router need not send anything more than a default route to the stub router.

Only specified routes are propagated from the stub router to the distribution router. As a stub router, the ASA responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” When the ASA is configured as a stub, it sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router depends on the distribution router to send the correct updates to all peers.

To enable the ASA as an EIGRP stub routing process, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 Click **Advanced** to configure the EIGRP stub routing process.

The Edit EIGRP Process Advanced Properties dialog box appears.

Step 5 In the Stub area on the Edit EIGRP Process Advanced Properties dialog box, choose one or more of the following EIGRP stub routing processes:

- 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.

Step 6 Click **OK**.

Step 7 Click the **Networks** tab.

Step 8 Click **Add** to add a new network entry.

The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 9 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 10 Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.

Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 11 Enter a network mask to apply to the IP address in the Network Mask field.

Step 12 Click **OK**.

Customizing EIGRP

This section describes how to customize the EIGRP routing and includes the following topics:

- Defining a Network for an EIGRP Routing Process, page 36-7
- Configuring Interfaces for EIGRP, page 36-7
- Configuring the Summary Aggregate Addresses on Interfaces, page 36-9
- Changing the Interface Delay Value, page 36-10
- Enabling EIGRP Authentication on an Interface, page 36-10
- Defining an EIGRP Neighbor, page 36-11
- Redistributing Routes Into EIGRP, page 36-12
- Filtering Networks in EIGRP, page 36-13
- Customizing the EIGRP Hello Interval and Hold Time, page 36-14
- Disabling Automatic Route Summarization, page 36-15
- Configuring Default Information in EIGRP, page 36-16
- Disabling EIGRP Split Horizon, page 36-17
- Restarting the EIGRP Process, page 36-17

Defining a Network for an EIGRP Routing Process

The Network table 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 configured for the EIGRP routing process. Each row of the table displays the network address and associated mask configured for the specified EIGRP routing process.

To add or define a network, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 Click the **Networks** tab.

Step 5 Click **Add** to add a new network entry.

The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 6 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 7 Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.

Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 8 Enter a network mask to apply to the IP address in the Network Mask field.

Step 9 Click **OK**.

Configuring Interfaces for EIGRP

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, you can configure the ASA that includes the network to which the interface is attached, and prevent that interface from sending or receiving EIGRP updates.

To configure interfaces for EIGRP, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane appears and displays the EIGRP interface configurations. The Interface Parameters table displays all of the interfaces on the ASA 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.

Step 5 Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The Edit EIGRP Interface Entry dialog box appears.

Step 6 In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 7 In the Hello Interval field, enter the interval between EIGRP hello packets sent on an interface.

Valid values range from 1 to 65535 seconds. The default value is 5 seconds.

Step 8 In the Hold Time field, enter the hold time, in seconds. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.

Step 9 Check the **Enable** check box for Split Horizon.

Step 10 In the Delay field, enter the delay value. The delay time is in tens of microseconds. Valid values range from 1 to 16777215.

Step 11 Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages.

Step 12 Enter the Key or Key ID values.

- In the Key field, enter the key to authenticate EIGRP updates. The key can contain up to 16 characters.
- In the Key ID field, enter the key identification value. Valid values range from 1 to 255.

Step 13 Click **OK**.

Configuring Passive Interfaces

You can configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates.

To configure passive interfaces, perform the following steps:

Note In ASDM, the Passive Interface table lists each interface that is configured as a passive interface.

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Click the **Passive Interfaces** tab.

Step 5 Choose the interface that you want to configure from the drop-down list.

Step 6 Check the **Suppress routing updates on all interfaces** check box to specify all interfaces as passive. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when the check box is checked.

Step 7 Click **Add** to add a passive interface entry.

The Add EIGRP Passive Interface dialog box appears. Choose the interface that you want to make passive and click **Add**. To remove a passive interface, choose the interface in the table and click **Delete**.

Step 8 Click **OK**.

Configuring the Summary Aggregate Addresses on Interfaces

You can configure a summary addresses on a per-interface basis. 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 an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

To create a summary address, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane shows the EIGRP interface configurations. The Interface Parameters table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see the "Configuring Interfaces for EIGRP" section on page 36-7.

Step 2 To configure the EIGRP parameters for an interface, double-click an interface entry or select the entry and click **Edit**.

Step 3 Click **OK**.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > 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.

Step 5 Click **Add** to add a new EIGRP summary address, or to click **Edit** to edit an existing EIGRP summary address in the table.

The Add Summary Address or Edit Summary Address dialog box appears. You can also double-click an entry in the table to edit that entry.

Step 6 In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 7 In the Interface drop-down list, choose the interface from which the summary address is advertised.

Step 8 In the IP Address field, enter the IP address of the summary route.

Step 9 In the Netmask field, choose or enter the network mask to apply to the IP address.

Step 10 Enter the administrative distance for the route in the Administrative Distance field. If left blank, the route has the default administrative distance of 5.

Step 11 Click **OK**.

Changing the Interface Delay Value

The interface delay value is used in EIGRP distance calculations. You can modify this value on a per-interface basis.

To change the interface delay value, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane shows the EIGRP interface configurations. The Interface Parameters table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see the "Configuring Interfaces for EIGRP" section on page 36-7.

Step 2 Double-click an interface entry or choose the Interface entry and click **Edit** to configure the delay value in the EIGRP parameters for an interface.

The Edit EIGRP Interface Entry dialog box appears.

Step 3 In the Delay field, enter the delay time, which is in tens of microseconds. Valid values are from 1 to 16777215.

Step 4 Click **OK**.

Enabling EIGRP Authentication on an Interface

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol. The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

EIGRP route authentication is configured on a per-interface basis. All EIGRP neighbors on interfaces configured for EIGRP message authentication must be configured with the same authentication mode and key for adjacencies to be established.

Note Before you can enable EIGRP route authentication, you must enable EIGRP.

To enable EIGRP authentication on an interface, perform the following steps:

Step 1 In the man ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Click the **Networks** tab.

Step 5 Click **Add** to add a new network entry.

The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 6 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 7 In the IP Address field, enter the IP address of the networks to participate in the EIGRP routing process.

Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 8 In the Network Mask field, choose or enter a network mask to apply to the IP address.

Step 9 Click **OK**.

Step 10 Choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane displays the EIGRP interface configurations. The Interface Parameters table displays all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see the "Configuring Interfaces for EIGRP" section on page 36-7.

Step 11 Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages. After you check this check box, provide one of the following:

- In the Key field, enter the key to authenticate EIGRP updates. The key can include up to 16 characters.
- In the Key ID field, enter the key identification value. Valid values range from 1 to 255.

Step 12 Click **OK**.

Defining an EIGRP Neighbor

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast 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.

To manually define an EIGRP neighbor, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Static Neighbor**.

The Static Neighbor pane appears and displays the statically-defined EIGRP neighbors. An EIGRP neighbor sends EIGRP routing information to and receives EIGRP routing information from the ASA. Normally, neighbors are dynamically discovered through the neighbor discovery process. However, on point-to-point, nonbroadcast 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.

From the Static Neighbor pane, you can add or edit a static neighbor.

Step 5 Click **Add** or **Edit** to add or edit a EIGRP static neighbor.

The Add or Edit EIGRP Neighbor Entry dialog box appears.

Step 6 Choose the EIGRP AS number from the drop-down list for the EIGRP process for which the neighbor is being configured.

Step 7 Choose the Interface Name from the Interface Name drop-down list, which is the interface through which the neighbor is available.

Step 8 Enter the IP address of the neighbor in the Neighbor IP Address field.

Step 9 Click **OK**.

Redistributing Routes Into EIGRP

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. You do not need to redistribute connected routes if they fall within the range of a **network** statement in the EIGRP configuration.

Note For RIP only: Before you begin this procedure, you must create a route map to further define which routes from the specified routing protocol are redistributed in to the RIP routing process. See Chapter 32, "Defining Route Maps," for more information about creating a route map.

To redistribute routes into the EIGRP routing process, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Redistribution**.

The Redistribution pane displays the rules for redistributing routes from other routing protocols to the EIGRP routing process. When redistributing static and connected routes to the EIGRP routing process, metrics are not required to be configured, although this is recommended. Each row of the Redistribution pane table includes a route redistribution entry.

Step 5 Click **Add** to add a new redistribution rule. If you are editing an existing redistribution rule, go to Step 6.

The Add EIGRP Redistribution Entry dialog box appears.

Step 6 Choose the address in the table and click **Edit** to edit an existing EIGRP static neighbor, You can also double-click an entry in the table to edit that entry.

The Edit EIGRP Redistribution Entry dialog box appears.

Step 7 Choose the AS number of the EIGRP routing process to which the entry applies from the drop-down list.

Step 8 In the Protocol area, click the radio button next to one of the following protocols for the routing process:

- **Static** to redistribute static routes to 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** to redistribute 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** to redistributes routes discovered by the RIP routing process to EIGRP.
- **OSPF** to redistribute routes discovered by the OSPF routing process to EIGRP.

Step 9 In the Optional Metrics area, choose one of the following metrics used for the redistributed route:

- **Bandwidth**, which is the EIGRP bandwidth metric in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay**, which is the EIGRP delay metric, in 10-microsecond units. Valid values range from 0 to 4294967295.
- **Reliability**, which is the EIGRP reliability metric. Valid values range from 0 to 255; 255 indicates 100 percent reliability.
- **Loading**, which is the EIGRP effective bandwidth (loading) metric. Valid values range from 1 to 255; 255 indicates 100 percent loaded.
- **MTU**, which is the MTU of the path. Valid values range from 1 to 65535.

Step 10 Choose the route map from the Route Map drop-down list to define which routes are redistributed into the EIGRP routing process. For more details about how to configure a route map, see Chapter 32, "Defining Route Maps."

Step 11 In the Optional OSPF Redistribution area, click one of the following OSPF radio buttons to further specify which OSPF routes are redistributed into the EIGRP routing process:

- **Match Internal** to match routes internal to the specified OSPF process.
- **Match External 1** to match type 1 routes external to the specified OSPF process.
- **Match External 2** to match type 2 routes external to the specified OSPF process.
- **Match NSSA-External 1** to match type 1 routes external to the specified OSPF NSSA.
- **Match NSSA-External 2** to match type 2 routes external to the specified OSPF NSSA.

Step 12 Click **OK**.

Filtering Networks in EIGRP

Note Before you begin this process, you must create a standard access list that defines the routes that you want to advertise. That is, create a standard access list that defines the routes that you want to filter from sending or receiving updates.

To filter networks in EIGRP, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Filter Rules**.

The Filter Rules pane appears and 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.

Step 5 Click **Add** to add a filter rule. If you are editing an already existing filter rule, skip to Step 6.

The Add Filter Rules dialog box appears.

Step 6 To edit a filter rule, choose the filter rule in the table and click **Edit**.

The Edit Filter Rules dialog appears. You can also double-click a filter rule to edit the rule. To remove a filter rule, choose the filter rule in the table and click **Delete**.

Step 7 Choose the AS number from the drop-down list of the EIGRP routing process to which the entry applies.

Step 8 Choose the direction of the filter routes from the drop-down list.

Choose **in** for rules that filter routes from incoming EIGRP routing updates. Choose **out** to filter routes from EIGRP routing updates that are sent by the ASA.

If you choose **out**, the Routing process field becomes active. Choose the type of route to be 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.

Step 9 Enter the OSPF process ID in the ID field.

Step 10 Click the **Interface** radio button and choose the interface to which the filter applies.

Step 11 Click **Add** or **Edit** to define an access list for the filter rule. Clicking **Edit** opens the Network Rule dialog box for the selected network rule.

The Network Rule dialog box appears.

Step 12 In the Action drop-down list, choose Permit to allow the specified network to be advertised; choose Deny to prevent the specified network from being advertised.

Step 13 In the IP Address field, 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**.

Step 14 From the Netmask drop-down list, choose 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.

Step 15 Click **OK**.

Customizing the EIGRP Hello Interval and Hold Time

The ASA periodically sends hello packets to discover neighbors and to learn when neighbors become unreachable or inoperative. By default, hello packets are sent every 5 seconds.

The hello packet advertises the ASA hold time. The hold time indicates to EIGRP neighbors the length of time the neighbor should consider the ASA reachable. If the neighbor does not receive a hello packet within the advertised hold time, then the ASA is considered unreachable. By default, the advertised hold time is 15 seconds (three times the hello interval).

Both the hello interval and the advertised hold time are configured on a per-interface basis. We recommend setting the hold time to be at minimum three times the hello interval.

To configure the hello interval and advertised hold time, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane appears and displays all of the EIGRP interface configurations.

Step 5 Double-click an interface entry or choose the entry and click **Edit**.

The Edit EIGRP Interface Entry dialog box appears.

Step 6 Choose the EIGRP AS number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.

Step 7 In the Hello Interval field, enter the interval between EIGRP hello packets sent on an interface.

Valid values range from 1 to 65535 seconds. The default value is 5 seconds.

Step 8 In the Hold Time field, specify the hold time, in seconds.

Valid values range from 1 to 65535 seconds. The default value is 15 seconds.

Step 9 Click **OK**.

Disabling Automatic Route Summarization

Automatic route summarization is enabled by default. The EIGRP routing process summarizes on network number boundaries. This can cause routing problems if you have noncontiguous networks.

For example, if you have a router with the networks 192.168.1.0, 192.168.2.0, and 192.168.3.0 connected to it, and those networks all participate in EIGRP, the EIGRP routing process creates the summary address 192.168.0.0 for those routes. If an additional router is added to the network with the networks 192.168.10.0 and 192.168.11.0, and those networks participate in EIGRP, they will also be summarized as 192.168.0.0. To prevent the possibility of traffic being routed to the wrong location, you should disable automatic route summarization on the routers creating the conflicting summary addresses.

To disable automatic route summarization in ASDM, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click the **Process Instance** tab.

Step 4 Click **Advanced**.

Step 5 In the Summary area, uncheck the **Auto-Summary** check box.

Note This setting is enabled by default.

Step 6 Click **OK**.

Configuring Default Information in EIGRP

You can control the sending and receiving of default route information in EIGRP updates. By default, default routes are sent and accepted. Configuring the ASA to disallow default information to be received causes the candidate default route bit to be blocked on received routes. Configuring the ASA to disallow default information to be sent disables the setting of the default route bit in advertised routes.

In ASDM, 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 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The main EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Do one of the following:

- Click **Add** to create a new entry.
- To edit an entry, double-click the entry in the table or select an entry in the table and click **Edit**.

  The Add Default Information or Edit Default Information dialog box appears for that entry. The EIGRP AS number is automatically selected in the EIGRP field.

Step 5 In the Direction field, choose the direction for the rule from the following options:

- **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 6 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, click **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, enter **0.0.0.0** as the network address and choose **0.0.0.0** as the network mask.

d. Click **OK** to add the specified network rule to the default information filter rule.

Step 7 Click **OK** to accept the default information filter rule.

Disabling EIGRP Split Horizon

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, update and query packets are not sent for destinations for which this interface is the next hop. Controlling update and query packets in this manner reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks, there may be situations where this behavior is not desired. For these situations, including networks in which you have EIGRP configured, you may want to disable split horizon.

If you disable split horizon on an interface, you must disable it for all routers and access servers on that interface.

To disable EIGRP split horizon, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane appears and displays the EIGRP interface configurations.

Step 2 Double-click an interface entry or choose the entry and click **Edit**.

The Edit EIGRP Interface Entry dialog box appears.

Step 3 Choose the EIGRP Autonomous system (AS) number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.

Step 4 Uncheck the **Split Horizon** check box.

Step 5 Click **OK**.

Restarting the EIGRP Process

To restart an EIGRP process or clear redistribution or counters, perform the following steps:

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Click **Reset**.

Monitoring EIGRP

You can use the following commands to monitor the EIGRP routing process. For examples and descriptions of the command output, see the command reference. Additionally, you can disable the logging of neighbor change messages and neighbor warning messages.

To monitor or disable various EIGRP routing statistics, perform the following steps:

Step 1 In the main ASDM window, choose **Monitoring > Routing > EIGRP Neighbor**.

Each row represents one EIGRP neighbor. For each neighbor, the list includes its IP address, the interface to which the neighbor is connected, the holdtime, the uptime, the queue length, the sequence number, the smoothed round trip time, and the retransmission timeout. The list of possible state changes are the following:

- NEW ADJACENCY—A new neighbor has been established.
- PEER RESTARTED—The other neighbor initiates the reset of the neighbor relationship. The router getting the message is not the one resetting the neighbor.
- HOLD TIME EXPIRED—The router has not heard any EIGRP packets from the neighbor within the hold-time limit.
- RETRY LIMIT EXCEEDED—EIGRP did not receive the acknowledgement from the neighbor for EIGRP reliable packets, and EIGRP has already tried to retransmit the reliable packet 16 times without any success.
- ROUTE FILTER CHANGED—The EIGRP neighbor is resetting because there is a change in the route filter.
- INTERFACE DELAY CHANGED—The EIGRP neighbor is resetting because there is a manual configuration change in the delay parameter on the interface.
- INTERFACE BANDWIDTH CHANGED—The EIGRP neighbor is resetting because there is a manual configuration change in the interface bandwidth on the interface.
- STUCK IN ACTIVE—The EIGRP neighbor is resetting because EIGRP is stuck in active state. The neighbor getting reset is the result of the stuck-in-active state.

Step 2 Click the EIGRP neighbor that you want to monitor.

Step 3 To remove the current list of neighbors, click **Clear Neighbors**.

Step 4 To refresh the current list of neighbors, click **Refresh**.

Note By default, neighbor change and neighbor warning messages are logged.

Feature History for EIGRP

Table 36-1 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 36-1 Feature History for EIGRP

Feature Name	Platform Releases	Feature Information
EIGRP Support	7.0(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP). We introduced the following screen: Configuration > Device Setup > Routing > EIGRP.
Dynamic Routing in Multiple Context Mode	9.0(1)	EIGRP routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > EIGRP > Setup.
Clustering	9.0(1)	For EIGRP, bulk synchronization, route synchronization, and layer 2 load balancing are supported in the clustering environment.

Configuring IPv6 Neighbor Discovery

This chapter describes how to enable and configure IPv6 neighbor discovery on the ASA and includes the following sections:

- Information About IPv6 Neighbor Discovery, page 37-1
- Licensing Requirements for IPv6 Neighbor Discovery, page 37-5
- Prerequisites for IPv6 Neighbor Discovery, page 37-5
- Guidelines and Limitations, page 37-5
- Default Settings for IPv6 Neighbor Discovery, page 37-7
- Configuring IPv6 Neighbor Discovery, page 37-7
- Viewing and Clearing Dynamically Discovered Neighbors, page 37-13
- Additional References, page 37-13
- Feature History for IPv6 Neighbor Discovery, page 37-14

Information About 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 readability of a neighbor, and keep track of neighboring routers.

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. In addition, 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.

This section includes the following topics:

- Neighbor Solicitation Messages, page 37-2
- Neighbor Reachable Time, page 37-3
- Duplicate Address Detection, page 37-3
- Router Advertisement Messages, page 37-3
- Static IPv6 Neighbors, page 37-5

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 37-1 shows the neighbor solicitation and response process.

Figure 37-1 IPv6 Neighbor Discovery—Neighbor Solicitation Message

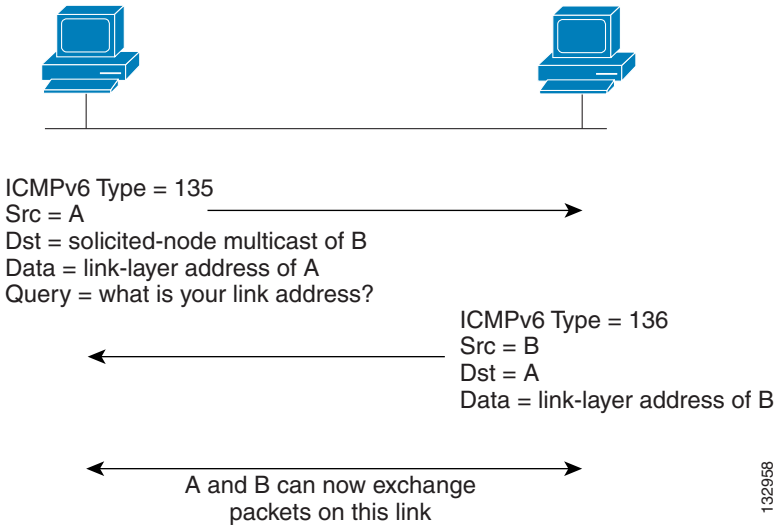

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

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.

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.

Duplicate Address Detection

During the stateless autoconfiguration process, Duplicate Address Detection 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:

```
%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 ASA uses neighbor solicitation messages to perform Duplicate Address Detection. By default, the number of times an interface performs Duplicate Address Detection is 1.

Router Advertisement Messages

An ASA can participate in router advertisements so that neighboring devices can dynamically learn a default router address. Router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface of the ASA. The router advertisement messages are sent to the all-nodes multicast address. Figure 37-2 shows an example of how to send router advertisement messages on an IPv6 configured interface.

Figure 37-2 IPv6 Neighbor Discovery—Router Advertisement Message

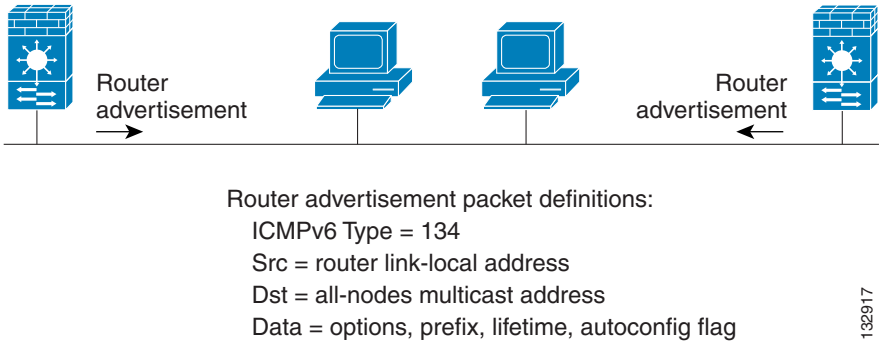

Router advertisement messages typically include the following information:

- 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 ASA to be the default router.
- The IPv6 network prefixes in use 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.

Static IPv6 Neighbors

You can manually define a neighbor in the IPv6 neighbor cache. 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. Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.

Licensing Requirements for IPv6 Neighbor Discovery

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites for IPv6 Neighbor Discovery

Configure IPv6 addressing according to the "Configuring IPv6 Addressing" section on page 17-15.

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 mode only. Transparent mode is not supported.

Additional Guidelines and Limitations

- The interval value is included in all IPv6 router advertisements that are sent out of this interface.
- The configured 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.
- The interval between transmissions should be less than or equal to the IPv6 router advertisement lifetime if the ASA is configured as a default router by using the **ipv6 nd ra-lifetime** command. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the specified value.
- The **ipv6 nd prefix** command allows control over the individual parameters per prefix, including whether or not the prefix should be advertised.

- By default, prefixes configured as addresses on an interface using the **ipv6 address** command are advertised in router advertisements. If you configure prefixes for advertisement using the **ipv6 nd prefix** command, then only these prefixes are advertised.
- The **default** keyword can be used to set default parameters for all prefixes.
- A date can be set to specify the expiration of a prefix. The valid and preferred lifetimes are counted down in real time. When the expiration date is reached, the prefix will no longer be advertised.
- When onlink is on (by default), the specified prefix is assigned to the link. Nodes sending traffic to such addresses that contain the specified prefix consider the destination to be locally reachable on the link.
- When autoconfig is on (by default), it indicates to hosts on the local link that the specified prefix can be used for IPv6 autoconfiguration.
- For stateless autoconfiguration to work correctly, the advertised prefix length in router advertisement messages must always be 64 bits.
- The router lifetime value is included in all IPv6 router advertisements sent out of the interface. The value indicates the usefulness of the ASA as a default router on this interface.
- Setting the value to a non-zero value indicates that the ASA should be considered a default router on this interface. The non-zero value for the router lifetime value should not be less than the router advertisement interval.

The following guidelines and limitations apply for configuring a static IPv6 neighbor:

- The **ipv6 neighbor** command is similar to the **arp** command. 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. These entries are stored in the configuration when the copy command is used to store the configuration.
- Use the **show ipv6 neighbor** command to view static entries in the IPv6 neighbor discovery cache.
- The **clear ipv6 neighbor** command deletes all entries in the IPv6 neighbor discovery cache except static entries. The **no ipv6 neighbor** command deletes a specified static entry from the neighbor discovery cache; the command does not remove dynamic entries—entries learned from the IPv6 neighbor discovery process—from the cache. Disabling IPv6 on an interface by using the **no ipv6 enable** command deletes all IPv6 neighbor discovery cache entries configured for that interface except static entries (the state of the entry changes to INCMP [Incomplete]).
- Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.
- The **clear ipv6 neighbor** command does not remove static entries from the IPv6 neighbor discovery cache; it only clears the dynamic entries.
- The ICMP syslogs generated are caused by a regular refresh of IPv6 neighbor entries. The ASA default timer for IPv6 neighbor entry is 30 seconds, so the ASA would generate ICMPv6 neighbor discovery and response packets about every 30 seconds. If the ASA has both failover LAN and state interfaces configured with IPv6 addresses, then every 30 seconds, ICMPv6 neighbor discovery and response packets will be generated by both ASAs for both configured and link-local IPv6 addresses. In addition, each packet will generate several syslogs (ICMP connection and local-host creation or teardown), so it may appear that constant ICMP syslogs are being generated. The refresh time for IPV6 neighbor entry is configurable on the regular data interface, but not configurable on the failover interface. However, the CPU impact for this ICMP neighbor discovery traffic is minimal.

Default Settings for IPv6 Neighbor Discovery

Table 37-1 lists the default settings for IPv6 Neighbor Discovery.

Table 37-1 Default IPv6 Neighbor Discovery Parameters

Parameters	Default
value for the neighbor solicitation transmission message interval	1000 seconds between neighbor solicitation transmissions.
value for the neighbor reachable time	The default is 0.
value for the router advertisement transmission interval	The default is 200 seconds.
value for the router lifetime	The default is 1800 seconds.
value for the number of consecutive neighbor solicitation messages sent during DAD	The default is one message.
prefix lifetime	The default lifetime is 2592000 seconds (30 days), and a preferred lifetime is 604800 seconds (7 days).
on-link flag	The flag is on by default, which means that the prefix is used on the advertising interface.
autoconfig flag	The flag is on by default, which means that the prefix is used for autoconfiguration.
static IPv6 neighbor	Static entries are not configured in the IPv6 neighbor discovery cache.

Configuring IPv6 Neighbor Discovery

- Configuring the Neighbor Solicitation Message Interval, page 37-8
- Configuring the Neighbor Reachable Time, page 37-8
- Configuring the Router Advertisement Transmission Interval, page 37-9
- Configuring the Router Lifetime Value, page 37-9
- Configuring DAD Settings, page 37-10
- Suppressing Router Advertisement Messages, page 37-10
- Configuring Address Config Flags for IPv6 DHCP Relay, page 37-11
- Configuring the IPv6 Prefix in Router Advertisements, page 37-11
- Configuring a Static IPv6 Neighbor, page 37-12

Configuring the Neighbor Solicitation Message Interval

To configure the interval between IPv6 neighbor solicitation retransmissions on an interface, perform the following steps.

Detailed 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 Addressing" section on page 17-15 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 running configuration.

Configuring the Neighbor Reachable Time

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.

Detailed 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 Addressing" section on page 17-15.

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 running configuration.

Configuring the Router Advertisement Transmission Interval

To configure the interval between IPv6 router advertisement transmissions on an interface, perform the following steps.

Detailed 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 Addressing" section on page 17-15.

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 running configuration.

Configuring the Router Lifetime Value

To configure the router lifetime value in IPv6 router advertisements on an interface, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration > Device Setup > Interfaces**.

Step 2 Select the interface you want to configure.

The interface must have been configured with an IPv6 address. For more information see the "Configuring IPv6 Addressing" section on page 17-15.

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 running configuration.

Configuring DAD Settings

To specify DAD settings on the interface, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration > Device Setup > Interfaces**.

Step 2 Select the interface you want to configure.

The interface must have been configured with an IPv6 address. For more information, see the "Configuring IPv6 Addressing" section on page 17-15.

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 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 range from 0 to 600. A zero value disables DAD processing on the specified interface. The default is one message.

Suppressing Router Advertisement Messages

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 ASA to supply the IPv6 prefix (for example, the outside interface).

To suppress the router lifetime value in IPv6 router advertisements on an interface, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration > Device Setup > Interfaces**.

Step 2 Select the interface for which you want to suppress the router advertisement transmissions. The interface must have been configured with an IPv6 address.

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.

Configuring Address Config Flags for IPv6 DHCP Relay

You can add a flag to IPv6 router advertisements to inform IPv6 autoconfigurtion clients to use DHCPv6 to obtain an IPv6 address and/or additional information such as the DNS server address.

Detailed Steps

Step 1 Choose **Configuration > Device Setup > Interfaces**.

Step 2 Select the interface that you want to configure.

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 **Hosts should use DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

Check the **Hosts should use DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Configuring the IPv6 Prefix in Router Advertisements

To configure the which IPv6 prefixes are included in IPv6 router advertisements, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration > Device Setup > Interfaces**.

Step 2 Select the interface for which you want to suppress the router advertisement transmissions. The interface must have been configured with an IPv6 address.

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 Interface IPv6 Prefixes area, click **Add**.

The Add IPv6 Prefix for Interface dialog box appears.

Step 6 Enter the IPv6 address with the prefix length.

Step 7 (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 8 (Optional) To indicate that the IPv6 prefix is not advertised, check the **No Advertisements** check box.

Step 9 (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 10 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 11 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 12 Click **OK** to save your settings.

The Interface IPv6 Prefixes Address field appears with the preferred and valid dates.

Configuring a Static IPv6 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 addresses, see the "Configuring IPv6 Addressing" section on page 17-15.

To add an IPv6 static neighbor, perform the following steps.

Detailed 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 running configuration.

Viewing and Clearing Dynamically Discovered 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 that neighbor.

To view dynamically discovered neighbors and clear these neighbors from the IPv6 neighbor discovery cache, perform the following steps:

Step 1 Choose **Monitoring > Interfaces > 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 dynamically discovered neighbor is removed from the cache.

Note This procedure clears only dynamically discovered neighbors from the cache; it does not clear static neighbors.

Additional References

For additional information related to implementing IPv6 prefixes, see the following topics:

- Related Documents for IPv6 Prefixes, page 37-14
- RFCs for IPv6 Prefixes and Documentation, page 37-14

Related Documents for IPv6 Prefixes

Related Topic	Document Title
ipv6 commands	*command reference*

RFCs for IPv6 Prefixes and Documentation

RFC	Title
RFC 2373 includes complete documentation to show how IPv6 network address numbers must be shown in router advertisements. The command argument *ipv6-prefix* indicates this network number, in which the address must be specified in hexadecimal format using 16-bit values between colons.	IP Version 6 Addressing Architecture
RFC 3849 specifies the requirements for using IPv6 address prefixes in documentation. The IPv6 unicast address prefix that has been reserved for use in documentation is 2001:DB8::/32.	IPv6 Address Prefix Reserved for Documentation

Feature History for IPv6 Neighbor Discovery

Table 37-2 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 37-2 Feature History for IPv6 Neighbor Discovery

Feature Name	Releases	Feature Information
IPv6 Neighbor Discovery	7.0(1)	We introduced this feature. We introduced the following screens: Monitoring > Interfaces > IPv6 Neighbor Discovery Cache. Configuration > Device Management > Advanced > IPv6 Neighbor Discovery Cache. Configuration > Device Setup > Interfaces > IPv6.
Address Config Flags for IPv6 DHCP Relay	9.0(1)	We modified the following screen: Configuration > Device Device Setup > Interfaces > IPv6.

PART 6

Configuring Network Address Translation

Information About NAT (ASA 8.3 and Later)

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

- Why Use NAT?, page 39-1
- NAT Terminology, page 39-2
- NAT Types, page 39-3
- NAT in Routed and Transparent Mode, page 39-12
- NAT and IPv6, page 39-15
- How NAT is Implemented, page 39-15
- NAT Rule Order, page 39-20
- Routing NAT Packets, page 39-21
- NAT for VPN, page 39-24
- DNS and NAT, page 39-30
- Where to Go Next, page 39-34

Note To start configuring NAT, see Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)," or Chapter 41, "Configuring Twice NAT (ASA 8.3 and Later)."

Why Use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.
- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only) (Version 9.0(1) and later)—If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.

Note NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT Terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the “real” network. Note that you can translate any network connected to the ASA, not just an inside network, Therefore if you configure NAT to translate outside addresses, “real” can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the “mapped” network.
- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address.

NAT Types

- NAT Types Overview, page 39-3
- Static NAT, page 39-3
- Dynamic NAT, page 39-8
- Dynamic PAT, page 39-10
- Identity NAT, page 39-12

NAT Types Overview

You can implement NAT using the following methods:

- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See the "Static NAT" section on page 39-3.
- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See the "Dynamic NAT" section on page 39-8.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See the "Dynamic PAT" section on page 39-10.
- Identity NAT—A real address is statically transalted to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See the "Identity NAT" section on page 39-12.

Static NAT

This section describes static NAT and includes the following topics:

- Information About Static NAT, page 39-3
- Information About Static NAT with Port Translation, page 39-4
- Information About One-to-Many Static NAT, page 39-6
- Information About Other Mapping Scenarios (Not Recommended), page 39-7

Information About Static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

Figure 39-1 shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 39-1 Static NAT

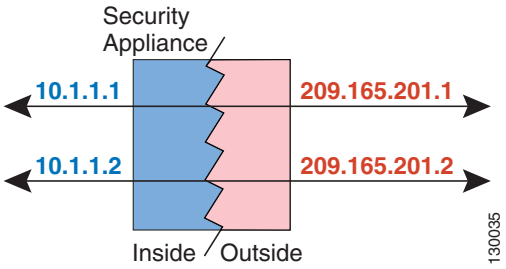

Note You can disable bidirectionality if desired.

Information About Static NAT with Port Translation

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

This section includes the following topics:

- Information About Static NAT with Port Address Translation, page 39-4
- Static NAT with Identity Port Translation, page 39-5
- Static NAT with Port Translation for Non-Standard Ports, page 39-5
- Static Interface NAT with Port Translation, page 39-6

Information About Static NAT with Port Address Translation

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

Figure 39-2 shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 39-2 Typical Static NAT with Port Translation Scenario

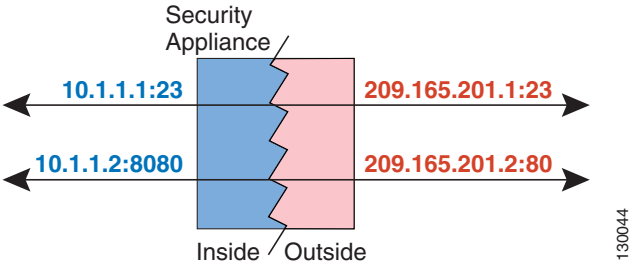

Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), the ASA automatically translates the secondary ports.

Static NAT with Identity Port Translation

The following static NAT with port translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT with port translation rules that use the same mapped IP address, but different ports. (See Figure 39-3. See the "Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation)" section on page 40-32 for details on how to configure this example.)

Figure 39-3 Static NAT with Port Translation

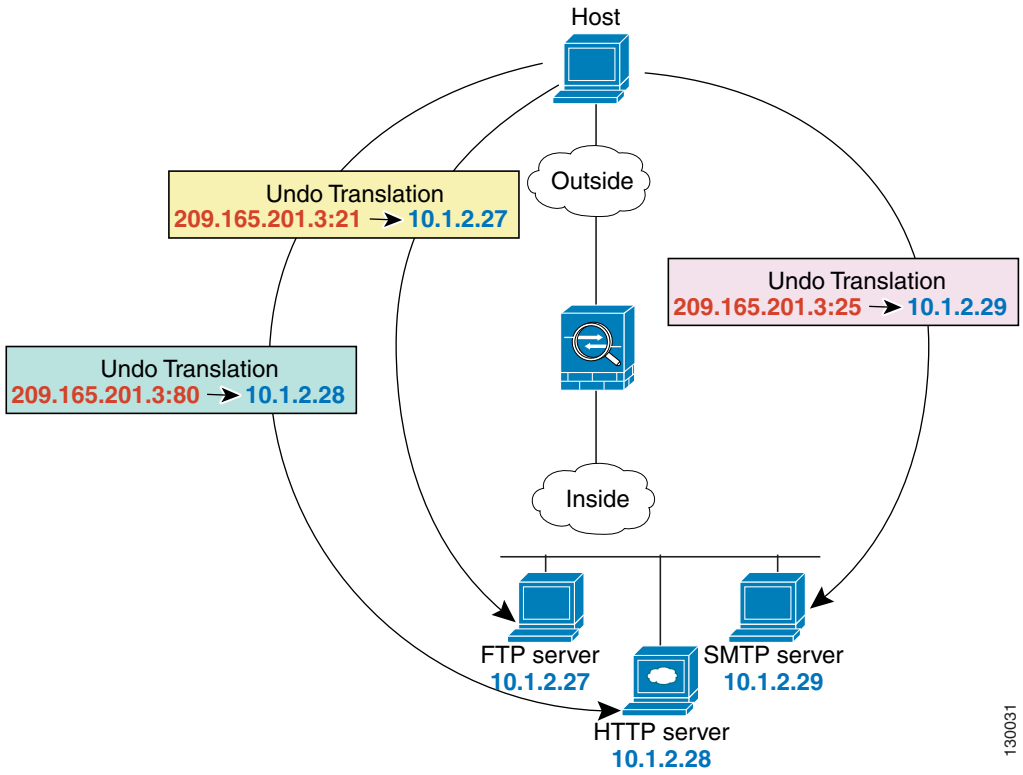

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the ASA outside interface to an inside host, then you can map the inside host IP address/port 23 to the ASA interface address/port 23. (Note that although Telnet to the ASA is not allowed to the lowest security interface, static NAT with interface port translation redirects the Telnet session instead of denying it).

Information About One-to-Many Static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

Figure 39-4 shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/1st mapped IP is technically the only bidirectional translation.

Figure 39-4 One-to-Many Static NAT

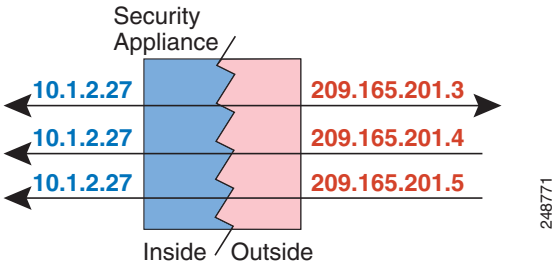

For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server (see Figure 39-5). (See the "Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many)" section on page 40-28 for details on how to configure this example.)

Figure 39-5 One-to-Many Static NAT

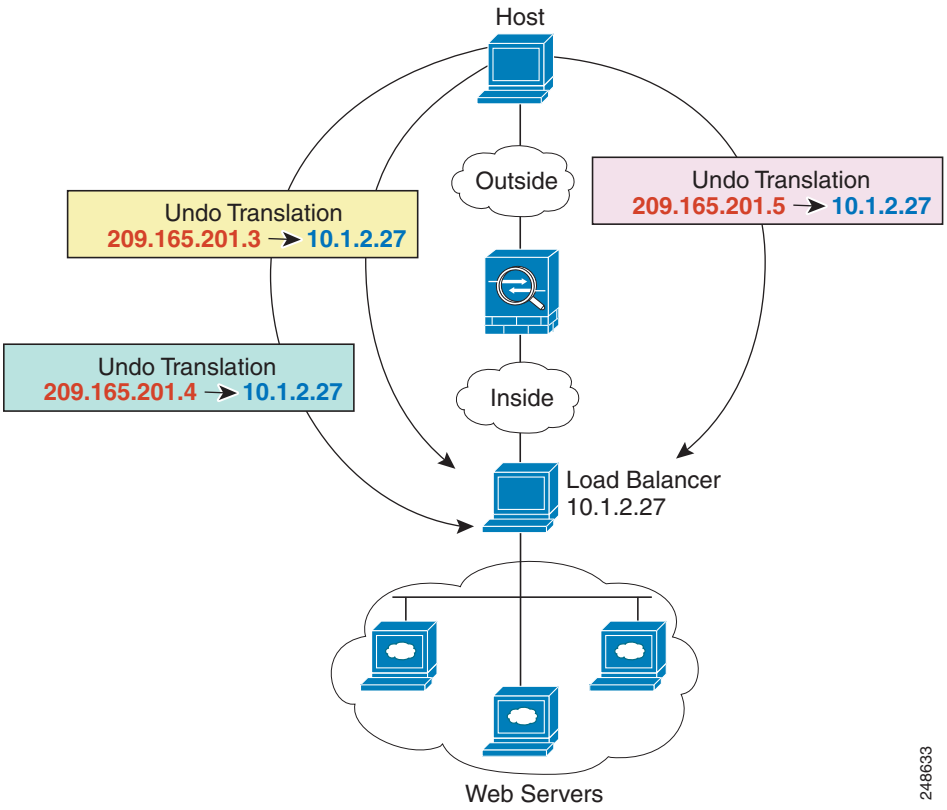

Information About Other Mapping Scenarios (Not Recommended)

The ASA has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

Figure 39-6 shows a typical few-to-many static NAT scenario.

Figure 39-6 Few-to-Many Static NAT

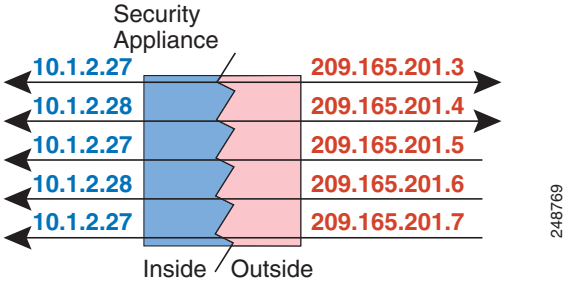

For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

Figure 39-7 shows a typical many-to-few static NAT scenario.

Figure 39-7 Many-to-Few Static NAT

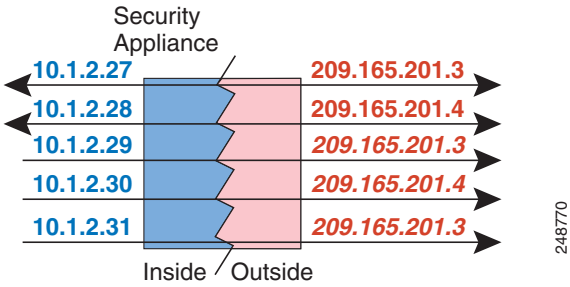

Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Dynamic NAT

This section describes dynamic NAT and includes the following topics:

- Information About Dynamic NAT, page 39-9
- Dynamic NAT Disadvantages and Advantages, page 39-10

Information About Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, the ASA assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Figure 39-8 shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 39-8 Dynamic NAT

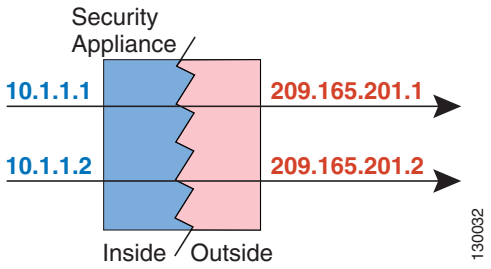

Figure 39-9 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the ASA drops the packet.

Figure 39-9 Remote Host Attempts to Initiate a Connection to a Mapped Address

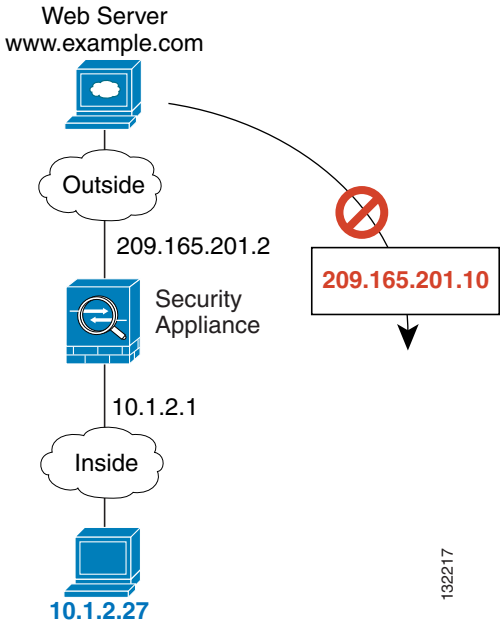

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

Dynamic NAT Disadvantages and Advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT or a PAT fallback method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

See the "Default Settings" section on page 57-4 for more information about NAT and PAT support.

Dynamic PAT

This section describes dynamic PAT and includes the following topics:

- Information About Dynamic PAT, page 39-10
- Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later), page 39-11
- Dynamic PAT Disadvantages and Advantages, page 39-11

Information About Dynamic PAT

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

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

Figure 39-10 shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 39-10 Dynamic PAT

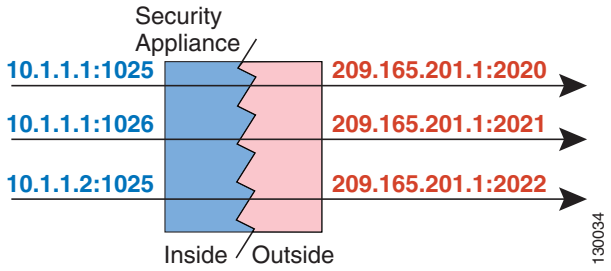

After the connection expires, the port translation also expires. For multi-session PAT, the PAT timeout is used, 30 seconds by default. For per-session PAT (9.0(1) and later), the xlate is immediately removed. Users on the destination network cannot reliably initiate a connection to a host that uses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

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

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

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

By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT be creating a per-session deny rule. See the "Configuring Per-Session PAT Rules" section on page 40-18.

Dynamic PAT Disadvantages and Advantages

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

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. See the "Default Settings" section on page 57-4 for more information about NAT and PAT support.

Dynamic PAT may also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. (8.4(2)/8.5(1) and later) You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate an address to itself. Identity NAT is necessary for remote access VPN, where you need to exempt the client traffic from NAT.

Figure 39-11 shows a typical identity NAT scenario.

Figure 39-11 Identity NAT

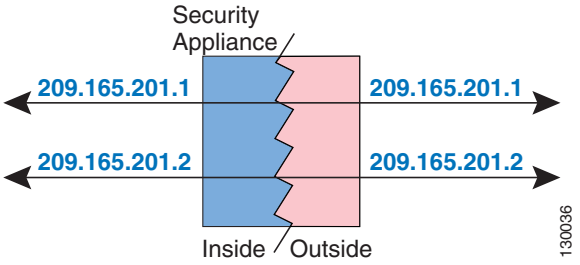

NAT in Routed and Transparent Mode

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

- NAT in Routed Mode, page 39-13
- NAT in Transparent Mode, page 39-13

NAT in Routed Mode

Figure 39-12 shows a typical NAT example in routed mode, with a private network on the inside.

Figure 39-12 NAT Example: Routed Mode

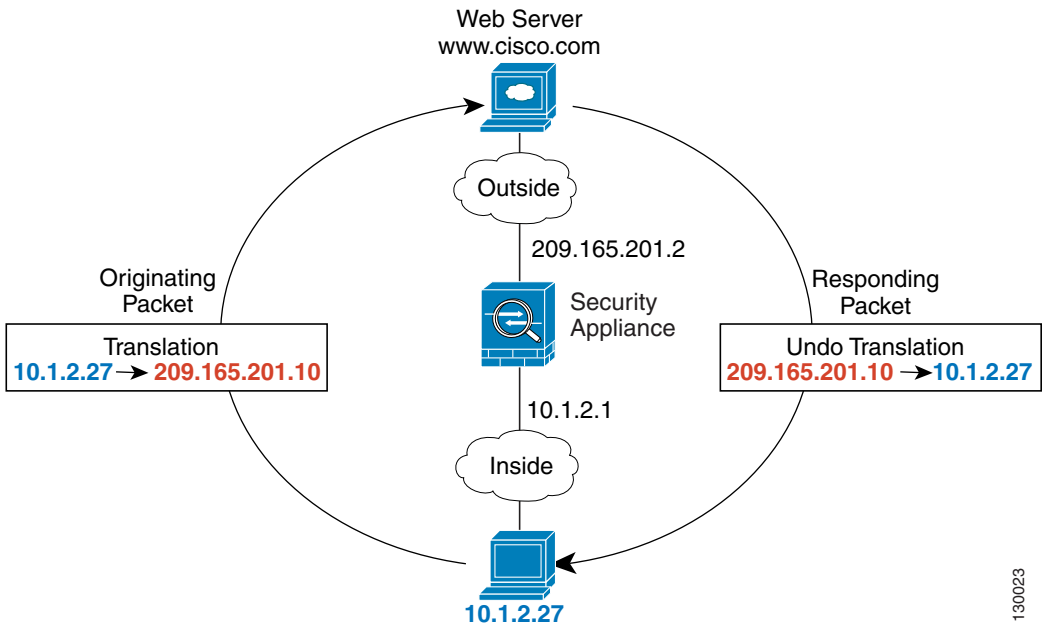

1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is changed to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the ASA receives the packet because the ASA performs proxy ARP to claim the packet.
3. The ASA then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in Transparent Mode

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

NAT in transparent mode has the following requirements and limitations:

- Because the transparent firewall does not have any interface IP addresses, you cannot use interface PAT.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the ASA sends an ARP request to a host on the other side of the ASA, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

Figure 39-13 shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 39-13 NAT Example: Transparent Mode

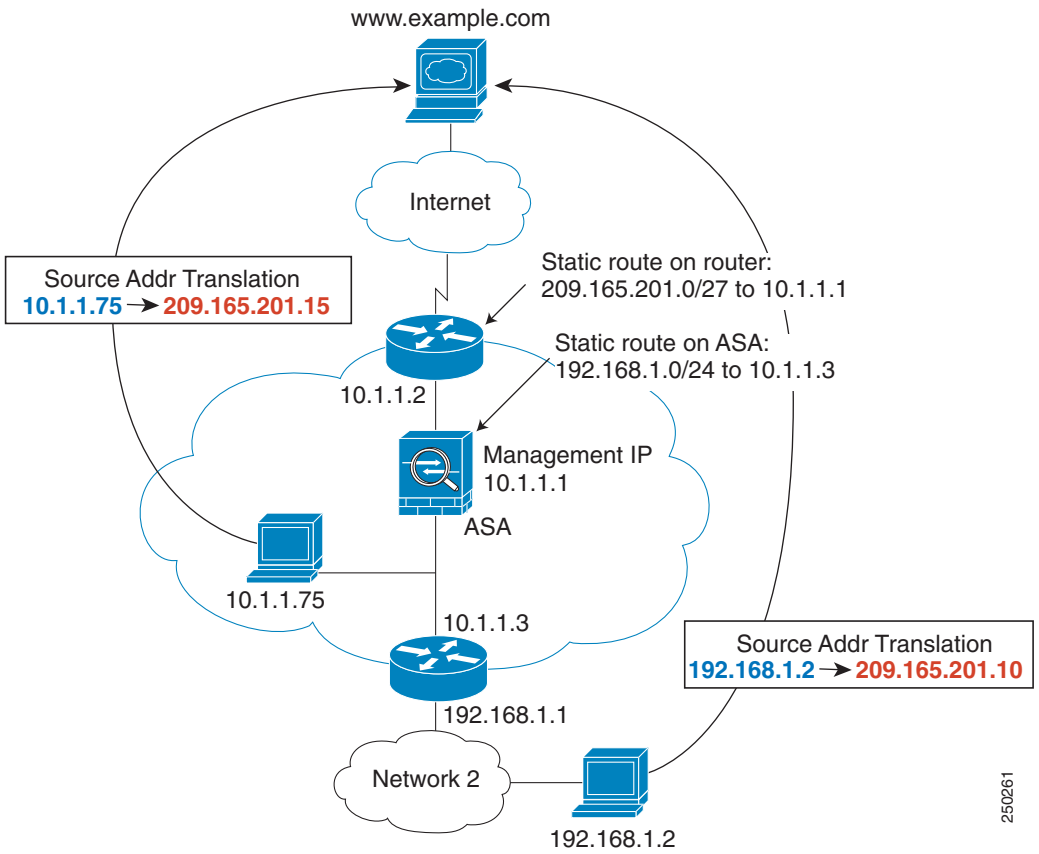

1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.
2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the ASA receives the packet because the upstream router includes this mapped network in a static route directed to the ASA management IP address. See the "Mapped Addresses and Routing" section on page 39-22 for more information about required routes.
3. The ASA then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the ASA sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the ASA looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the ASA static route for 192.168.1.0/24. See the "Transparent Mode Routing Requirements for Remote Networks" section on page 39-24 for more information about required routes.

NAT and IPv6

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (twice NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-tonet, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

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

How NAT is Implemented

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

- Main Differences Between Network Object NAT and Twice NAT, page 39-15
- Information About Network Object NAT, page 39-16
- Information About Twice NAT, page 39-16

Main Differences Between Network Object NAT and Twice NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Network object NAT—You define NAT as a parameter for a network object. A network object names an IP host, range, or subnet so you can then use the object in configuration instead of the actual IP addresses. The network object IP address serves as the real address. This method lets you easily add NAT to network objects that might already be used in other parts of your configuration.
  - Twice NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that twice NAT is more scalable.
- How source and destination NAT is implemented.

  - Network object NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
  - Twice NAT—A single rule translates both the source and destination. A matching packet only matches the one rule, and further rules are not checked. Even if you do not configure the optional destination address for twice NAT, a matching packet still only matches one twice NAT rule. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.

- Order of NAT Rules.
  - Network object NAT—Automatically ordered in the NAT table.
  - Twice NAT—Manually ordered in the NAT table (before or after network object NAT rules).

  See the "NAT Rule Order" section on page 39-20 for more information.

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

Information About Network Object NAT

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

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

When a packet enters the ASA, both the source and destination IP addresses are checked against the network object NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use twice NAT for that kind of functionality (twice NAT lets you identify the source and destination address in a single rule).

To start configuring network object NAT, see Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)."

Information About Twice NAT

Twice NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

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

To start configuring twice NAT, see Chapter 41, "Configuring Twice NAT (ASA 8.3 and Later)."

Figure 39-14 shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130. (See the "Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation)" section on page 40-32 for details on how to configure this example.)

Figure 39-14 Twice NAT with Different Destination Addresses

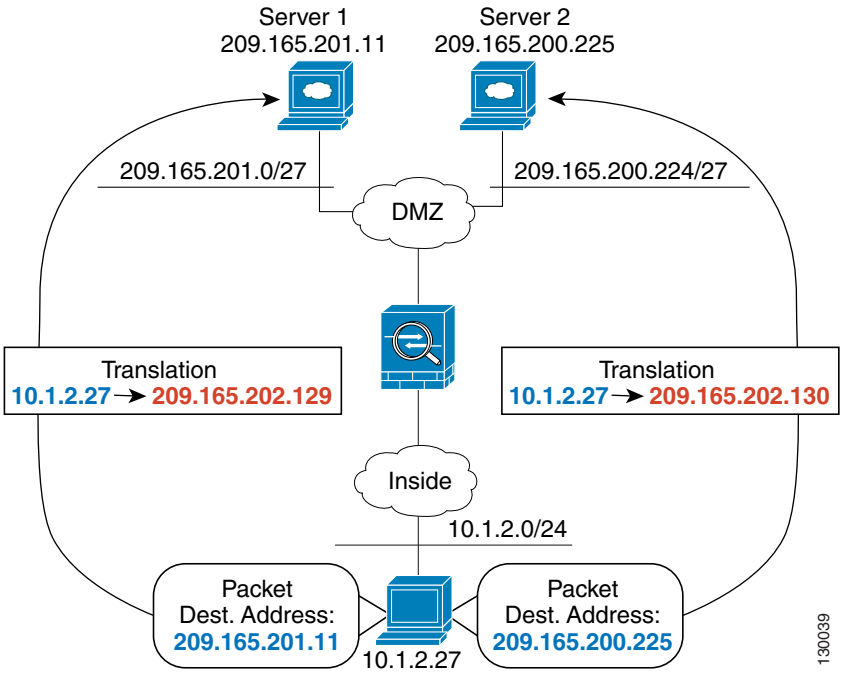

Figure 39-15 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 39-15 Twice NAT with Different Destination Ports

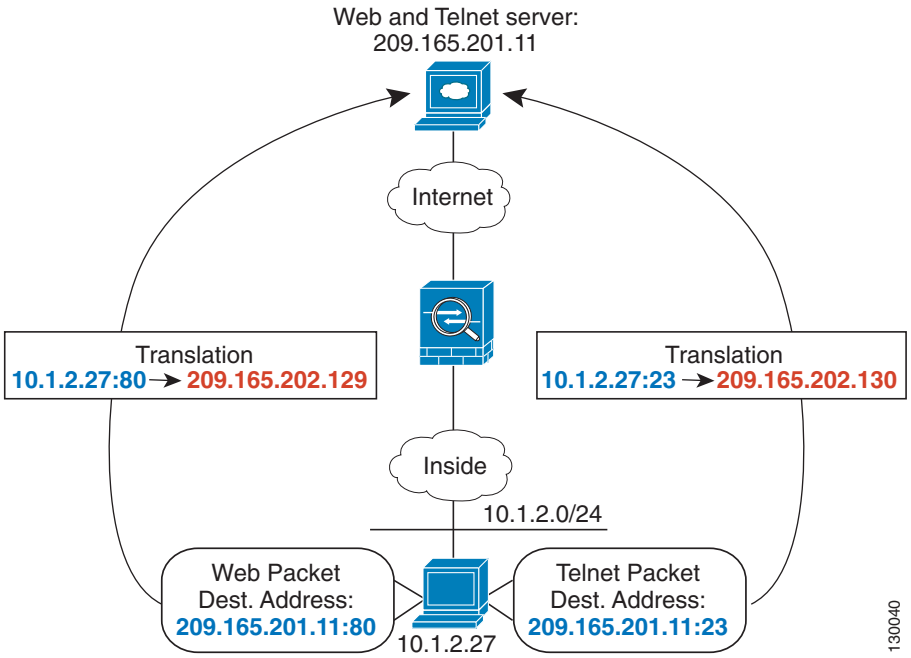

Figure 39-16 shows a remote host connecting to a mapped host. The mapped host has a twice static NAT translation that translates the real address only for traffic to and from the 209.165.201.0/27 network. A translation does not exist for the 209.165.200.224/27 network, so the translated host cannot connect to that network, nor can a host on that network connect to the translated host.

Figure 39-16 Twice Static NAT with Destination Address Translation

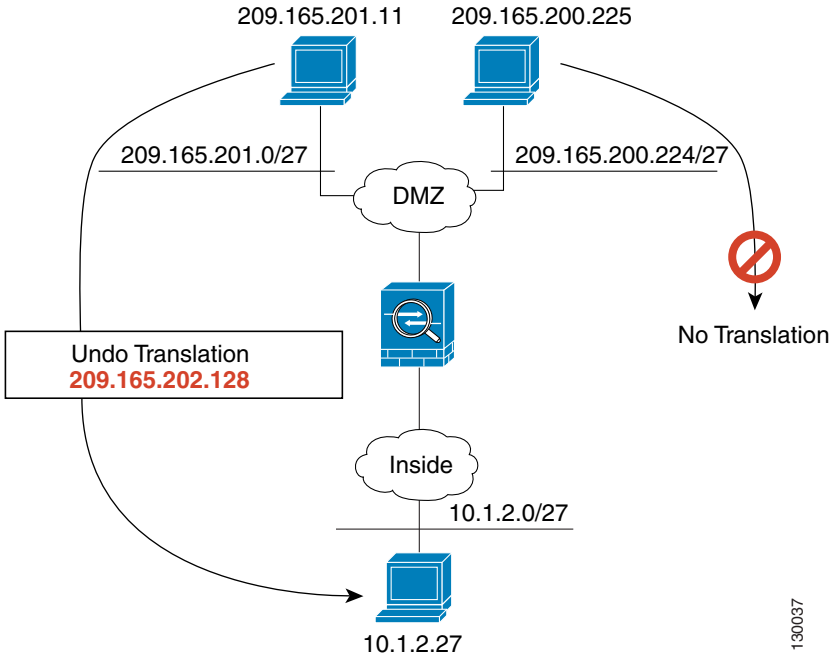

NAT Rule Order

Network object NAT rules and twice NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3. Table 39-1 shows the order of rules within each section.

Table 39-1 NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Twice NAT	Applied on a first match basis, in the order they appear in the configuration. By default, twice NAT rules are added to section 1. **Note** If you configure EasyVPN remote, the ASA dynamically adds invisible NAT rules to the end of this section. Be sure that you do not configure a twice NAT rule in this section that might match your VPN traffic, instead of matching the invisible rule. If VPN does not work due to NAT failure, consider adding twice NAT rules to section 3 instead.
Section 2	Network object NAT	Section 2 rules are applied in the following order, as automatically determined by the ASA: 1. Static rules. 2. Dynamic rules. Within each rule type, the following ordering guidelines are used: a. Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. b. For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. c. If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Twice NAT	Section 3 rules are applied on a first match basis, in the order they appear in the configuration. You can specify whether to add a twice NAT rule to section 3 when you add the rule.

For section 2 rules, for example, you have the following IP addresses defined within network objects:

192.168.1.0/24 (static)

192.168.1.0/24 (dynamic)

10.1.1.0/24 (static)

192.168.1.1/32 (static)

172.16.1.0/24 (dynamic) (object def)

172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:

192.168.1.1/32 (static)

10.1.1.0/24 (static)

192.168.1.0/24 (static)

172.16.1.0/24 (dynamic) (object abc)

172.16.1.0/24 (dynamic) (object def)

192.168.1.0/24 (dynamic)

NAT Interfaces

You can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside (Figure 39-17).

Figure 39-17 Specifying Any Interface

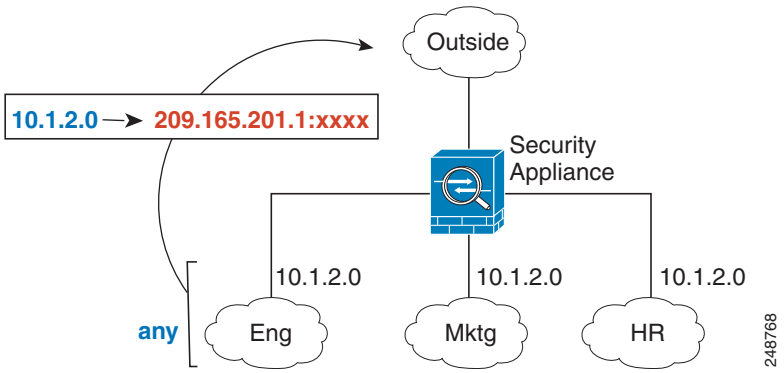

Note For transparent mode, you must choose specific source and destination interfaces.

Routing NAT Packets

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

- Mapped Addresses and Routing, page 39-22
- Transparent Mode Routing Requirements for Remote Networks, page 39-24
- Determining the Egress Interface, page 39-24

Mapped Addresses and Routing

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

See additional guidelines about mapped IP addresses in Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)," and Chapter 41, "Configuring Twice NAT (ASA 8.3 and Later)."

See the following mapped address types:

- Addresses on the same network as the mapped interface.

  If you use addresses on the same network as the mapped interface, the ASA uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the ASA does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.

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

- Addresses on a unique network.

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

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

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

  (8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. **Note**: You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

  Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for "any" IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly-connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches "any" address). The ASA will then proxy ARP for the address, even though the packet is not actually destined for the ASA. (Note that this problem occurs even if you have a twice NAT rule; although

the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the "source" address). If the ASA ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the ASA (see Figure 39-18).

Figure 39-18 Proxy ARP Problems with Identity NAT

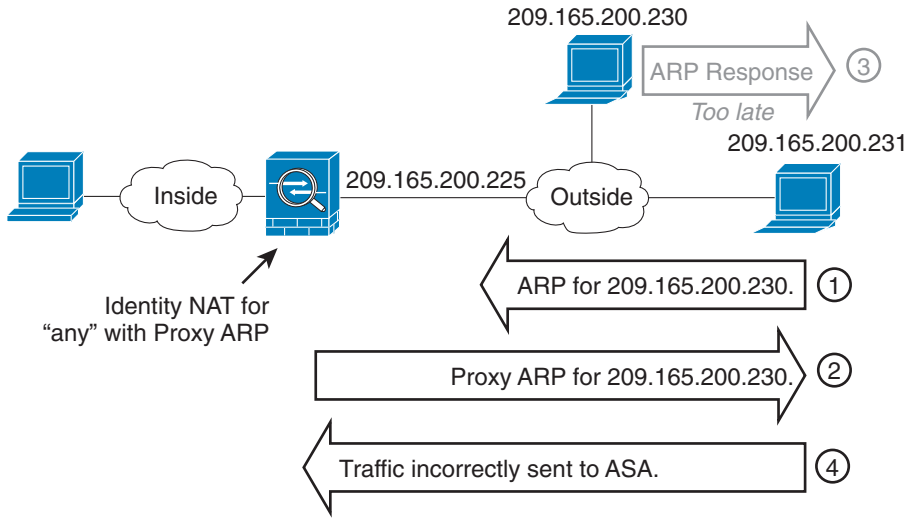

In rare cases, you need proxy ARP for identity NAT; for example for virtual Telnet. When using AAA for network access, a host needs to authenticate with the ASA using a service like Telnet before any other traffic can pass. You can configure a virtual Telnet server on the ASA to provide the necessary login. When accessing the virtual Telnet address from the outside, you must configure an identity NAT rule for the address specifically for the proxy ARP functionality. Due to internal processes for virtual Telnet, proxy ARP lets the ASA keep traffic destined for the virtual Telnet address rather than send the traffic out the source interface according to the NAT rule. (See Figure 39-19).

Figure 39-19 Proxy ARP and Virtual Telnet

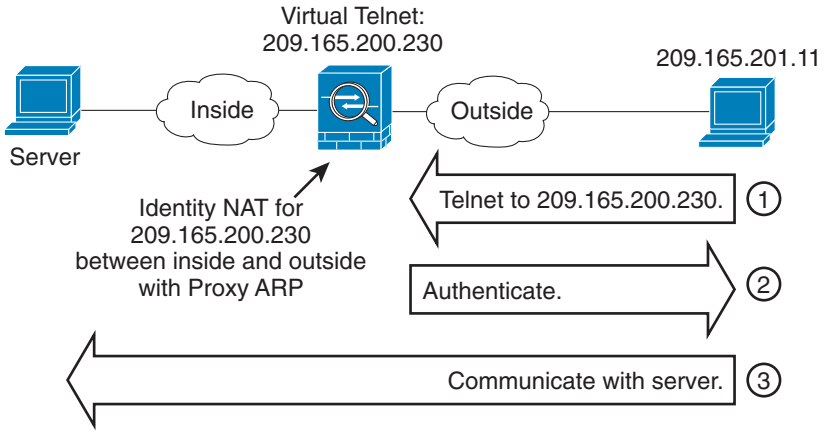

Transparent Mode Routing Requirements for Remote Networks

If the ASA performs NAT for a host that is not on the directly-connected network, then you need to configure a static route on the ASA for that network. You also need to have a static route for embedded IP addresses that are at least one hop away from the ASA (such as in VoIP or DNS traffic) when you have inspection and NAT enabled.

Determining the Egress Interface

In transparent mode, the ASA determines the egress interface for a NAT packet by using the NAT configuration; you must specify the source and destination interfaces as part of the NAT configuration.

In routed mode, the ASA determines the egress interface for a NAT packet in the following way:

- If you specify an optional interface, then the ASA uses the NAT configuration to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration, but you have the option to always use a route lookup instead.
- If you do not specify a specific interface, then the ASA uses a route lookup to determine the egress interface.

NAT for VPN

- NAT and Remote Access VPN, page 39-24
- NAT and Site-to-Site VPN, page 39-26
- NAT and VPN Management Access, page 39-28
- Troubleshooting NAT and VPN, page 39-30

NAT and Remote Access VPN

Figure 39-20 shows both an inside server (10.1.1.6) and a VPN client (209.165.201.10) accessing the Internet. Unless you configure split tunnelling for the VPN client (where only specified traffic goes through the VPN tunnel), then Internet-bound VPN traffic must also go through the ASA. When the VPN traffic enters the ASA, the ASA decrypts the packet; the resulting packet includes the VPN client local address (10.3.3.10) as the source. For both inside and VPN client local networks, you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. To allow the VPN traffic to exit the same interface it entered, you also need to enable intra-interface communication (AKA “hairpin” networking).

Figure 39-20 Interface PAT for Internet-Bound VPN Traffic (Intra-Interface)

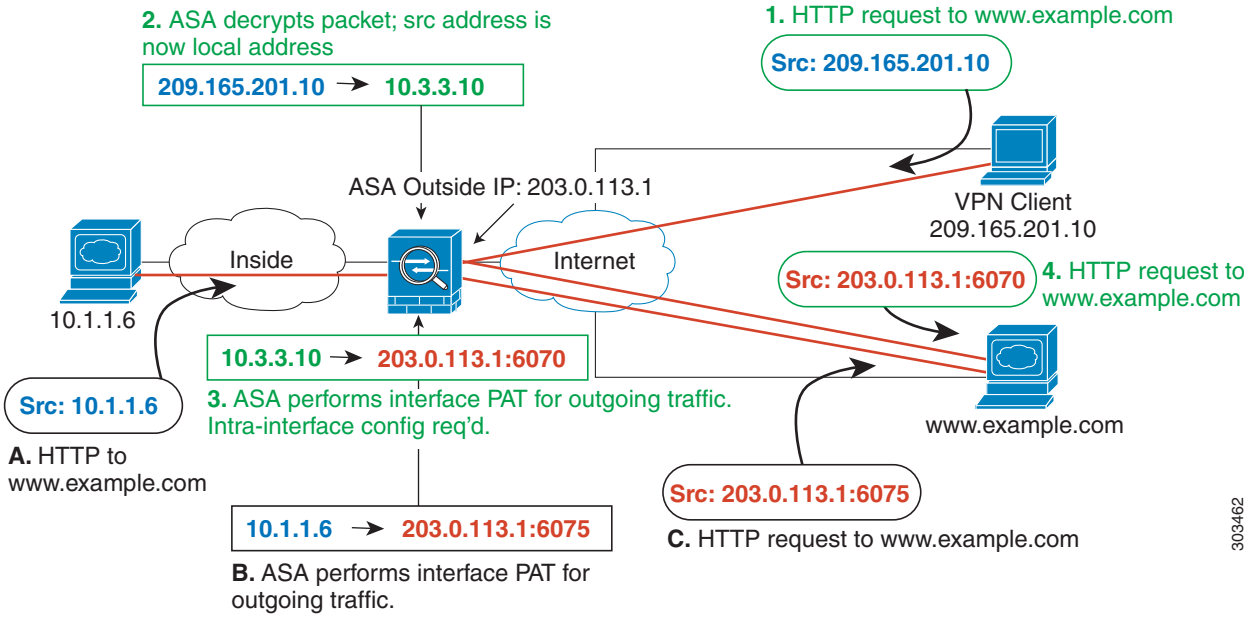

Figure 39-21 shows a VPN client that wants to access an inside mail server. Because the ASA expects traffic between the inside network and any outside network to match the interface PAT rule you set up for Internet access, traffic from the VPN client (10.3.3.10) to the SMTP server (10.1.1.6) will be dropped due to a reverse path failure: traffic from 10.3.3.10 to 10.1.1.6 does not match a NAT rule, but returning traffic from 10.1.1.6 to 10.3.3.10 *should* match the interface PAT rule for outgoing traffic. Because forward and reverse flows do not match, the ASA drops the packet when it is received. To avoid this failure, you need to exempt the inside-to-VPN client traffic from the interface PAT rule by using an identity NAT rule between those networks. Identity NAT simply translates an address to the same address.

Figure 39-21 Identity NAT for VPN Clients

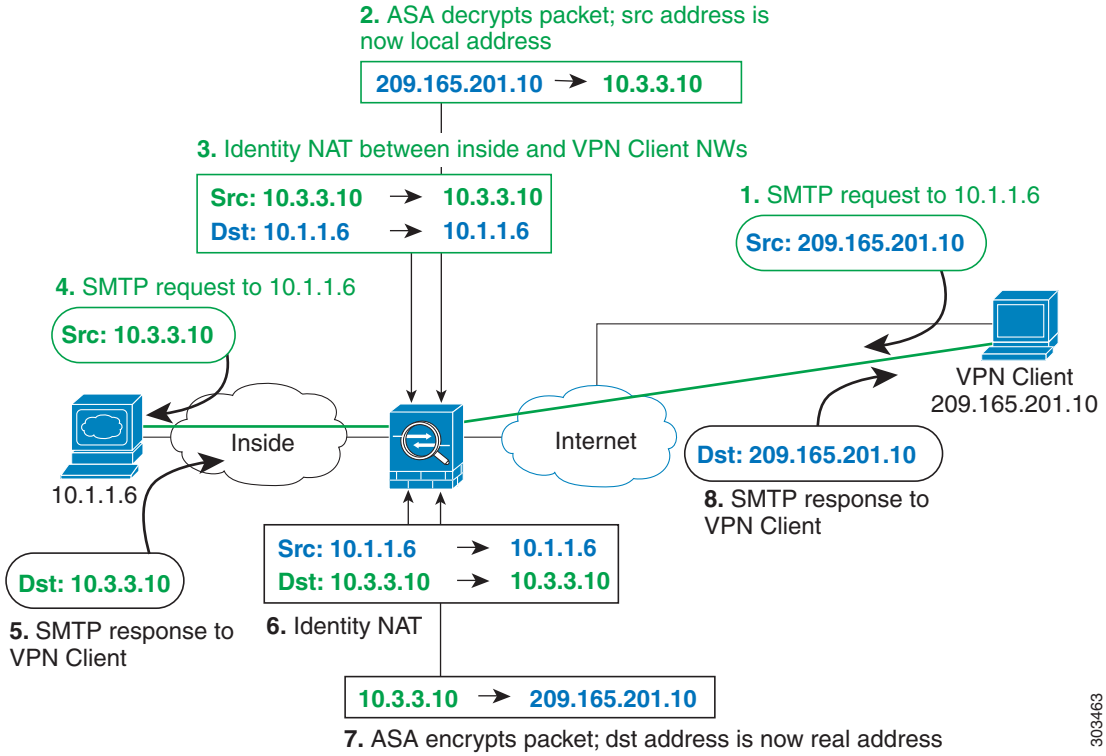

See the following sample NAT configuration for the above network:

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

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

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

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

NAT and Site-to-Site VPN

Figure 39-22 shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 39-22 Interface PAT and Identity NAT for Site-to-Site VPN

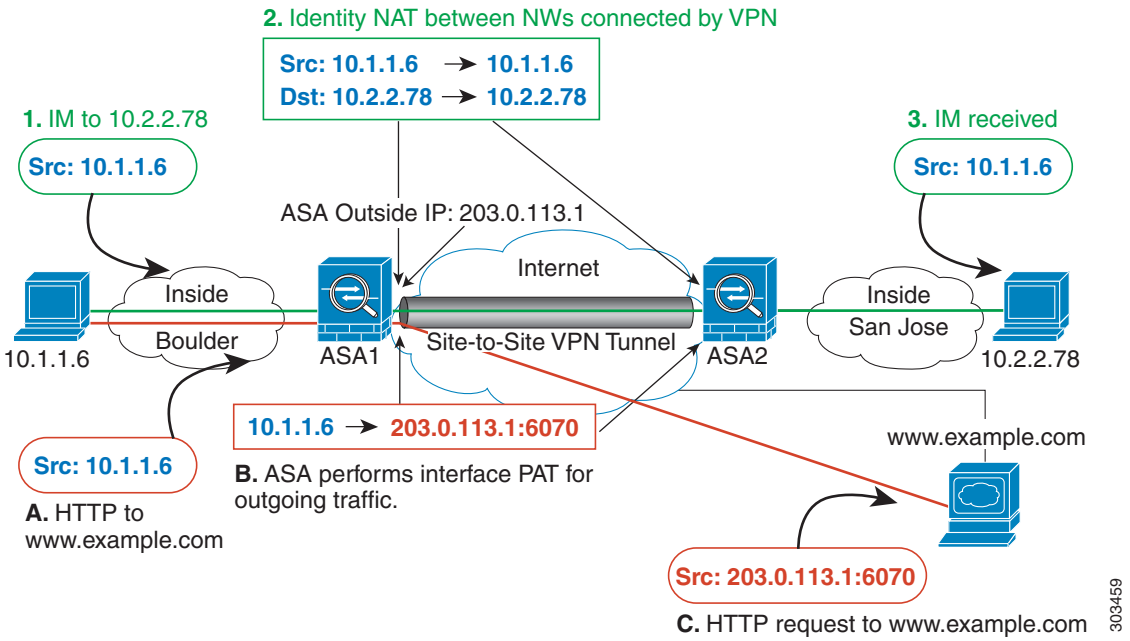

Figure 39-23 shows a VPN client connected to ASA1 (Boulder), with a Telnet request for a server (10.2.2.78) accessible over a site-to-site tunnel between ASA1 and ASA2 (San Jose). Because this is a hairpin connection, you need to enable intra-interface communication, which is also required for non-split-tunneled Internet-bound traffic from the VPN client. You also need to configure identity NAT between the VPN client and the Boulder & San Jose networks, just as you would between any networks connected by VPN to exempt this traffic from outbound NAT rules.

Figure 39-23 VPN Client Access to Site-to-Site VPN

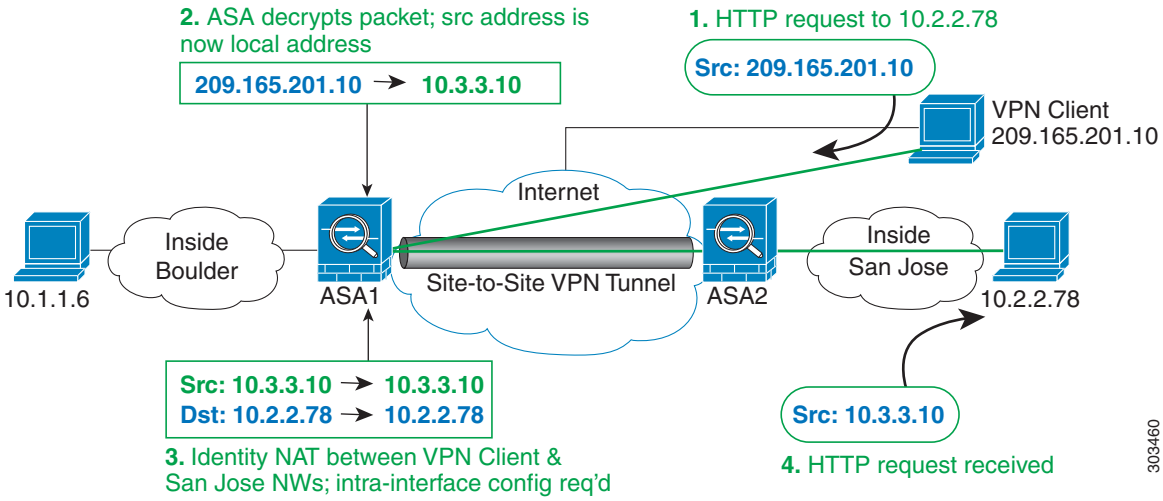

See the following sample NAT configuration for ASA1 (Boulder):

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NAT and VPN Management Access

When using VPN, you can allow management access to an interface other than the one from which you entered the ASA ("Configuring Management Access Over a VPN Tunnel" section on page 34-18). For example, if you enter the ASA from the outside interface, the management-access feature lets you connect to the inside interface using ASDM, SSH, Telnet, or SNMP; or you can ping the inside interface.

Figure 39-24 shows a VPN client Telnetting to the ASA inside interface. When you use a management-access interface, and you configure identity NAT according to the "NAT and Remote Access VPN" or "NAT and Site-to-Site VPN" section, you must configure NAT with the route lookup option. Without route lookup, the ASA sends traffic out the interface specified in the NAT command, regardless of what the routing table says; in the below example, the egress interface is the inside interface. You do not want the ASA to send the management traffic out to the inside network; it will never return to the inside interface IP address. The route lookup option lets the ASA send the traffic directly to the inside interface IP address instead of to the inside network. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected. See the "Determining the Egress Interface" section on page 39-24 for more information about the route lookup option.

Figure 39-24 VPN Management Access

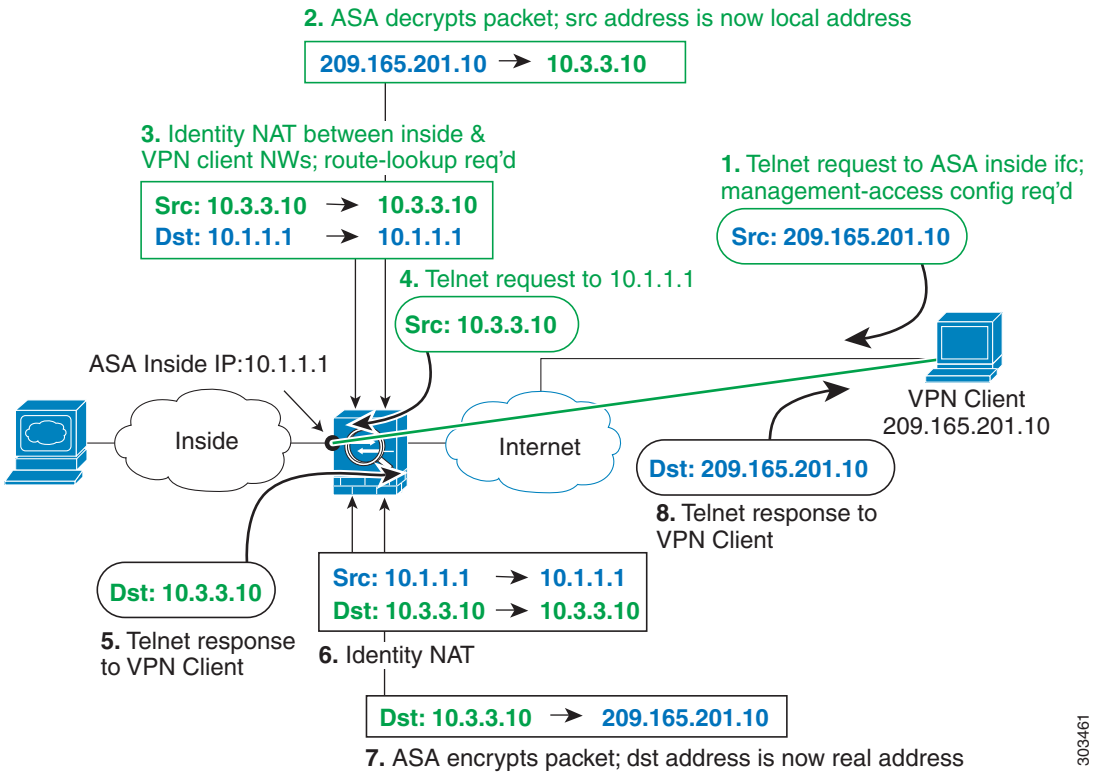

See the following sample NAT configuration for the above network:

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

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

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

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

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

Troubleshooting NAT and VPN

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

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

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

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

DNS and NAT

You might need to configure the ASA to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule.

For IPv4, this feature rewrites the A record, or address record, in DNS replies that match a NAT rule. For IPv6, it rewrites the AAAA record. For NAT64, it converts between A and AAAA records. For DNS replies traversing from a mapped interface to any other interface, the address record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the address record is rewritten from the real value to the mapped value.

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

Figure 39-25 shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network. In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address. When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The ASA refers to the static rule for the inside server

and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

Figure 39-25 DNS Reply Modification, DNS Server on Outside

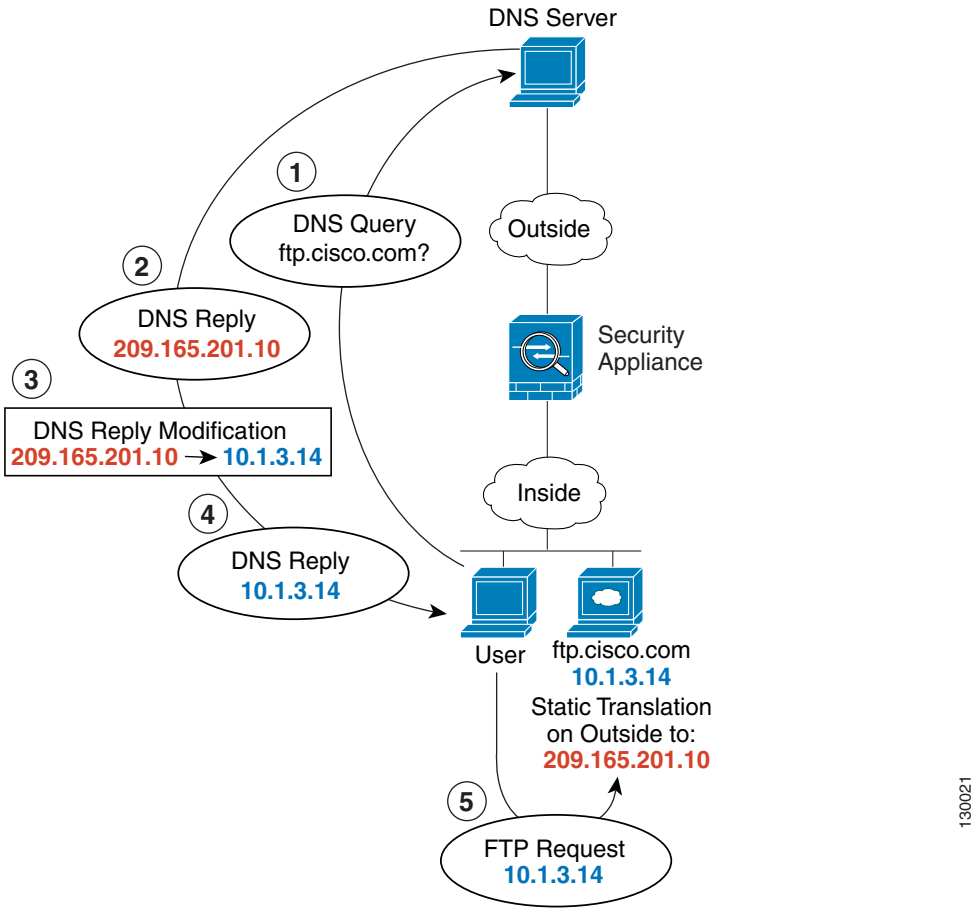

Figure 39-26 shows a user on the inside network requesting the IP address for ftp.cisco.com, which is on the DMZ network, from an outside DNS server. The DNS server replies with the mapped address (209.165.201.10) according to the static rule between outside and DMZ even though the user is not on the DMZ network. The ASA translates the address inside the DNS reply to 10.1.3.14. If the user needs to access ftp.cisco.com using the real address, then no further configuration is required. If there is also

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

Figure 39-26 DNS Reply Modification, DNS Server, Host, and Server on Separate Networks

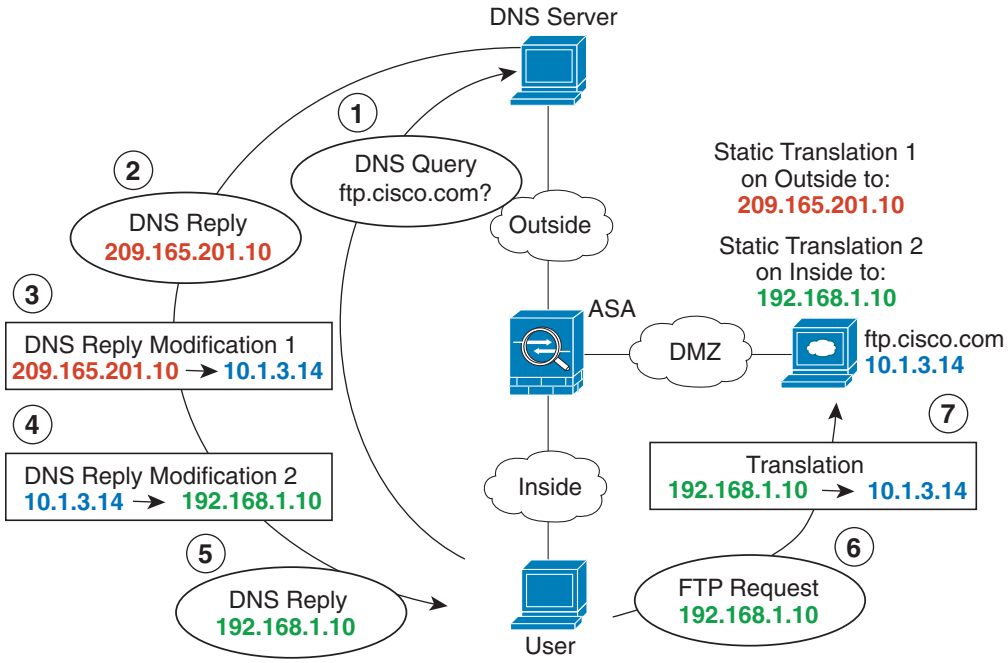

Figure 39-27 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 39-27 DNS Reply Modification, DNS Server on Host Network

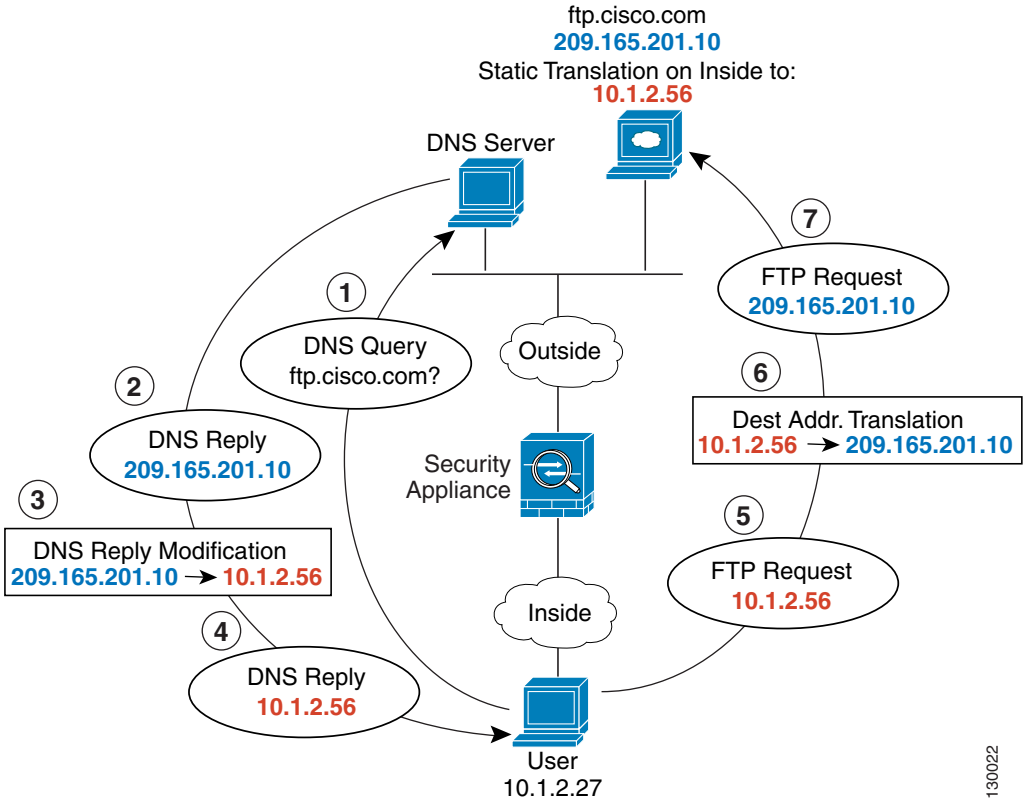

Figure 39-27 shows an FTP server and DNS server on the outside IPv4 network. The ASA has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.

Figure 39-28 DNS64 Reply Modification Using Outside NAT

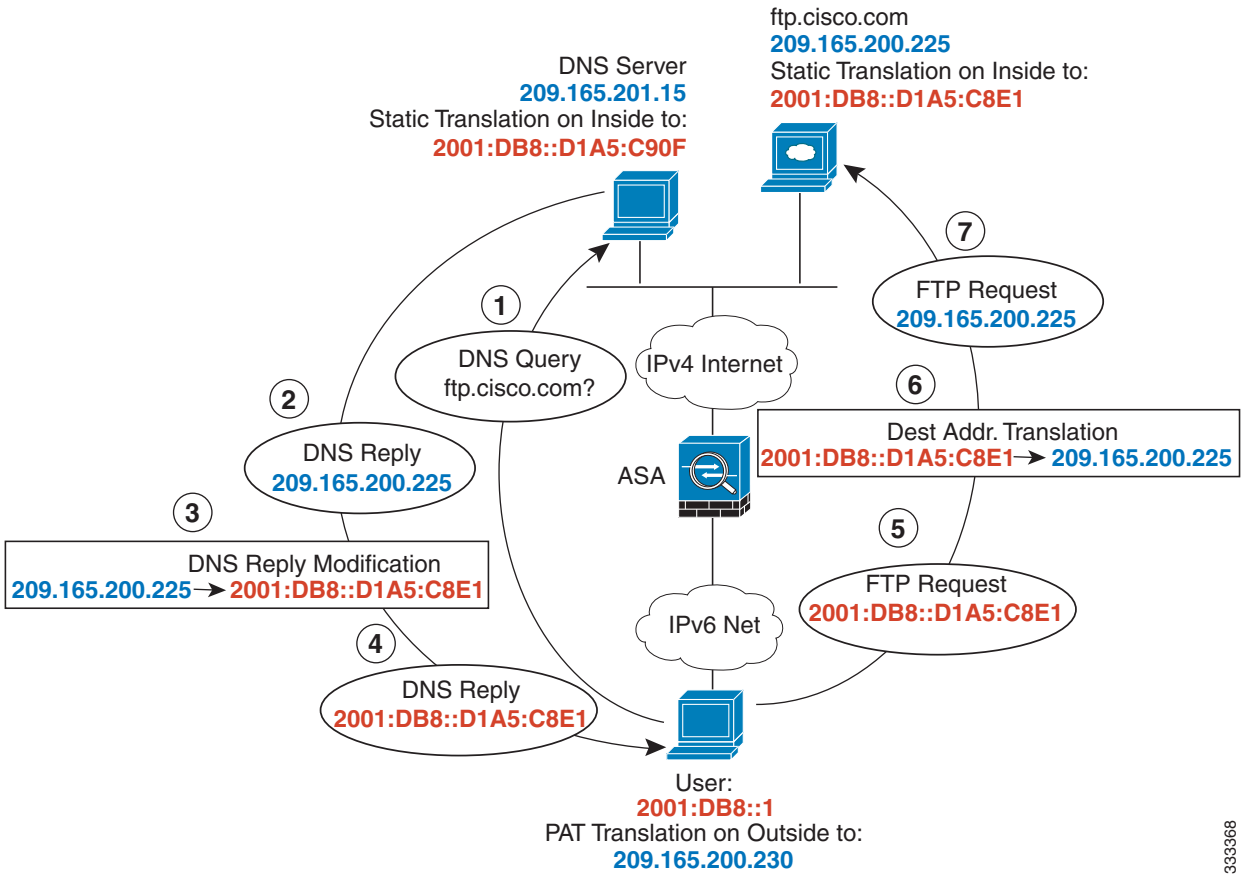

Where to Go Next

To configure network object NAT, see Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)."

To configure twice NAT, see Chapter 41, "Configuring Twice NAT (ASA 8.3 and Later)."

Configuring Network Object NAT (ASA 8.3 and Later)

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

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

- Information About Network Object NAT, page 40-1
- Licensing Requirements for Network Object NAT, page 40-2
- Prerequisites for Network Object NAT, page 40-2
- Guidelines and Limitations, page 40-2
- Default Settings, page 40-3
- Configuring Network Object NAT, page 40-4
- Monitoring Network Object NAT, page 40-19
- Configuration Examples for Network Object NAT, page 40-20
- Feature History for Network Object NAT, page 40-45

Note For detailed information about how NAT works, see Chapter 39, "Information About NAT (ASA 8.3 and Later)."

Information About Network Object NAT

When a packet enters the ASA, both the source and destination IP addresses are checked against the network object NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that a source address should be translated to A when going to destination X, but be translated to B when going to destination Y. Use twice NAT for that kind of functionality (twice NAT lets you identify the source and destination address in a single rule).

For detailed information about the differences between twice NAT and network object NAT, see the "How NAT is Implemented" section on page 39-15.

Network object NAT rules are added to section 2 of the NAT rules table. For more information about NAT ordering, see the "NAT Rule Order" section on page 39-20.

Licensing Requirements for Network Object NAT

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites for Network Object NAT

Depending on the configuration, you can configure the mapped address inline if desired or you can create a separate network object or network object group for the mapped address. Network object groups are particularly useful for creating a mapped address pool with discontinous IP address ranges or multiple hosts or subnets. To create a network object or group, see the "Configuring Network Objects and Groups" section on page 25-2.

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

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

- Supported in routed and transparent firewall mode.
- In transparent mode, you must specify the real and mapped interfaces; you cannot use --Any--.
- In transparent mode, you cannot configure interface PAT, because the transparent mode interfaces do not have IP addresses. You also cannot use the management IP address as a mapped address.
- In transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

IPv6 Guidelines

- Supports IPv6. See also the "NAT and IPv6" section on page 39-15.
- For routed mode, you can also translate between IPv4 and IPv6.
- For transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.
- For transparent mode, a PAT pool is not supported for IPv6.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.

- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

Additional Guidelines

- You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address, for example, **object network obj-10.10.10.1-01**, **object network obj-10.10.10.1-02**, and so on.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command. However, clearing the translation table disconnects all current connections that use translations.

Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** command. This safeguard ensures that the same address is not assigned to multiple hosts.

- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.
- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify --Any-- interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), use the interface name instead of the IP address.
  - (Transparent mode) The management IP address.
  - (Dynamic NAT) The standby interface IP address when VPN is enabled.
  - Existing VPN pool addresses.
- For application inspection limitations with NAT or PAT, see the "Default Settings" section on page 57-4 in Chapter 57, "Getting Started with Application Layer Protocol Inspection."

Default Settings

- (Routed mode) The default real and mapped interface is Any, which applies the rule to all interfaces.
- (8.3(1), 8.3(2), and 8.4(1)) The default behavior for identity NAT has proxy ARP disabled. You cannot configure this setting. (8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. See the "Routing NAT Packets" section on page 39-21 for more information.
- If you specify an optional interface, then the ASA uses the NAT configuration to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration, but you have the option to always use a route lookup instead. See the "Routing NAT Packets" section on page 39-21 for more information.

Configuring Network Object NAT

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

- Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool, page 40-4
- Configuring Dynamic PAT (Hide), page 40-8
- Configuring Static NAT or Static NAT-with-Port-Translation, page 40-11
- Configuring Identity NAT, page 40-15
- Configuring Per-Session PAT Rules, page 40-18

Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool

This section describes how to configure network object NAT for dynamic NAT or for dynamic PAT using a PAT pool. For more information, see the "Dynamic NAT" section on page 39-8 or the "Dynamic PAT" section on page 39-10.

Guidelines

For a PAT pool:

- If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you have a lot of traffic that uses the lower port ranges, you can now specify for a PAT pool a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535.
- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT and a flat range, then the other rule must also specify extended PAT and a flat range.

For extended PAT for a PAT pool:

- Many application inspections do not support extended PAT. See the "Default Settings" section on page 57-4 in Chapter 57, "Getting Started with Application Layer Protocol Inspection," for a complete list of unsupported inspections.
- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.
- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.

For round robin for a PAT pool:

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

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

Detailed Steps

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

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

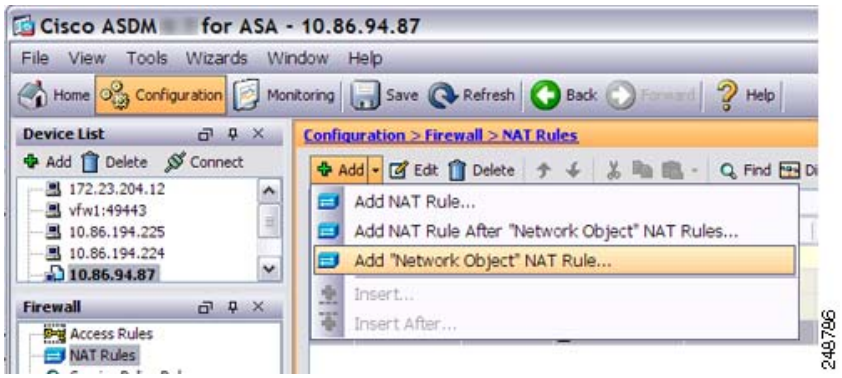

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the "Configuring a Network Object" section on page 25-3.

The Add/Edit Network Object dialog box appears.

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

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

b. Type—Host, Network, or Range.

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

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

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

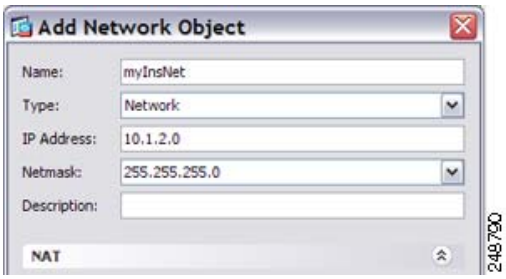

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

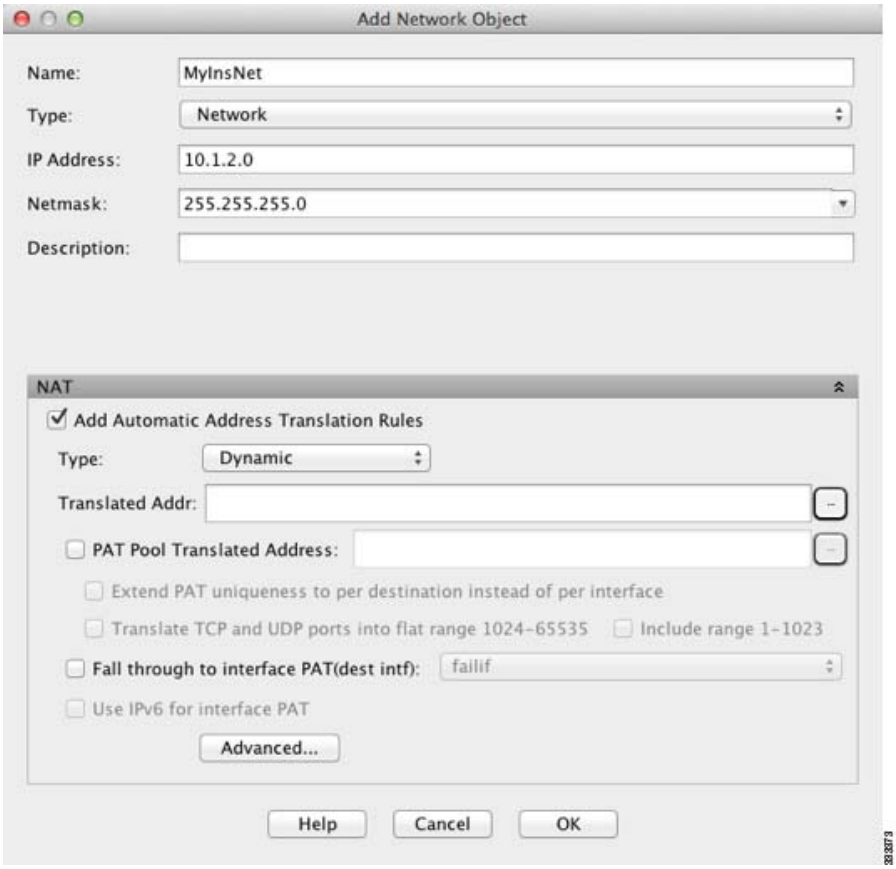

Step 4 Check the **Add Automatic Translation Rules** check box.

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

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

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

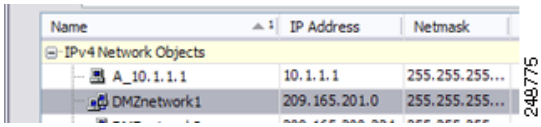

Note The object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

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

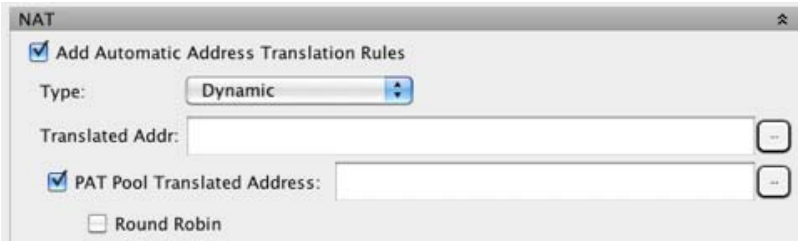

a. Do not enter a value for the Translated Addr. field; leave it blank.

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

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

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

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

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

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

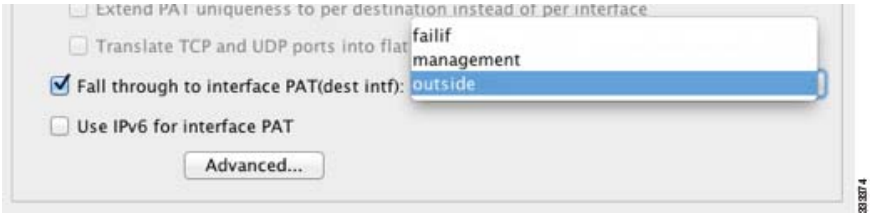

Step 8 (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

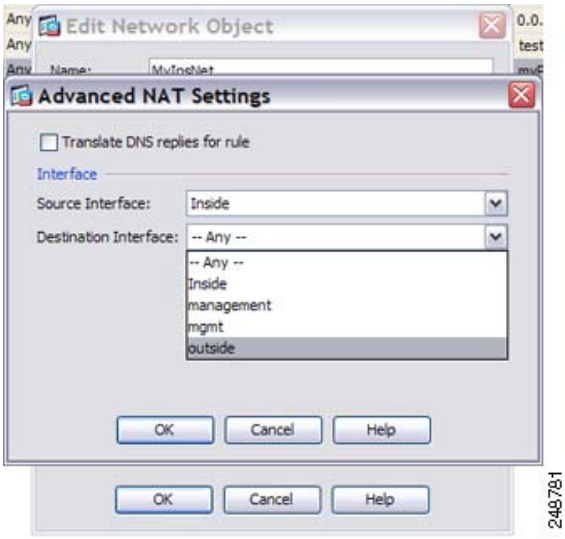

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See the "DNS and NAT" section on page 39-30 for more information.
- (Required for Transparent Firewall Mode) Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
- (Required for Transparent Firewall Mode) Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

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

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

Configuring Dynamic PAT (Hide)

This section describes how to configure network object NAT for dynamic PAT (hide). For dynamic PAT using a PAT pool, see the "Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool" section on page 40-4 instead of using this section. For more information, see the "Dynamic PAT" section on page 39-10.

Detailed Steps

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

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

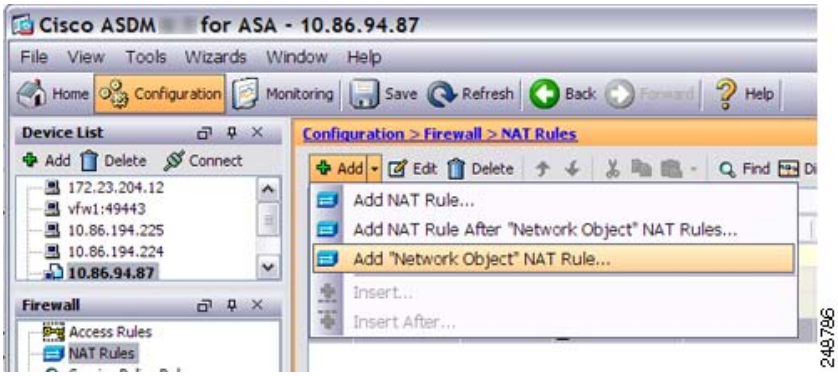

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the "Configuring a Network Object" section on page 25-3.

The Add/Edit Network Object dialog box appears.

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

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

b. Type—Host, Network, or Range.

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

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

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

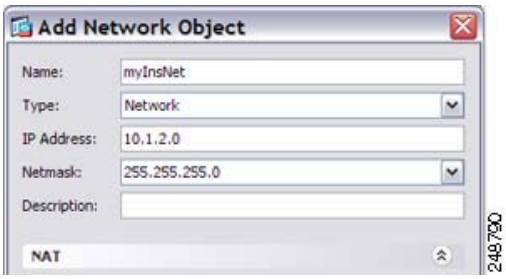

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

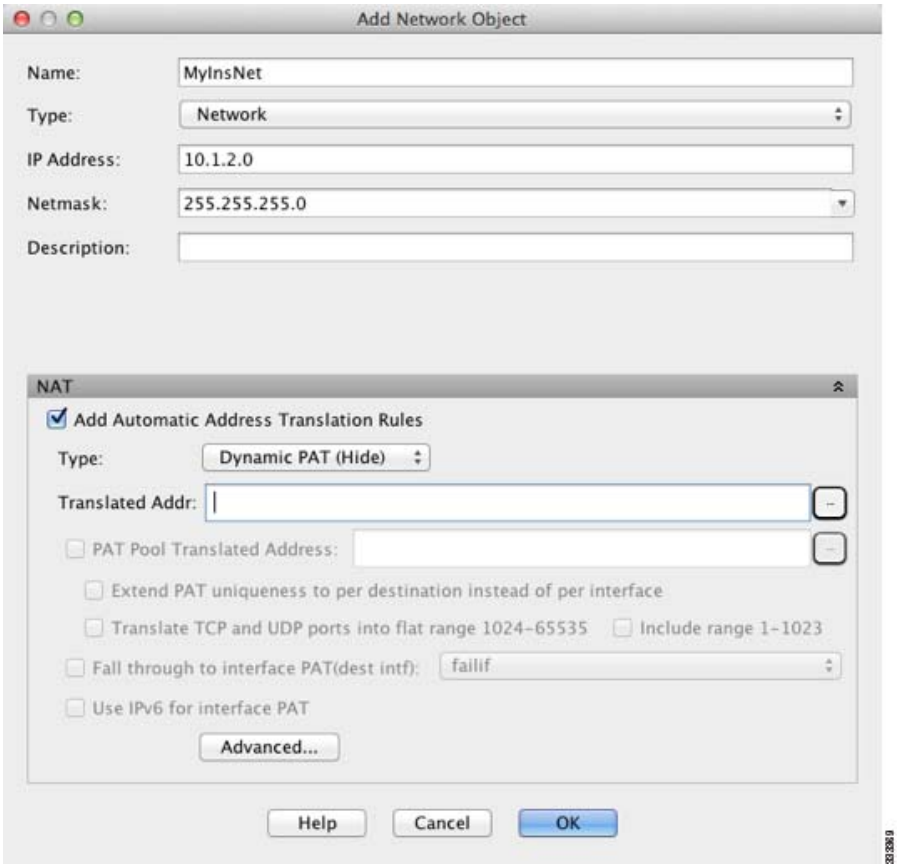

Step 4 Check the **Add Automatic Translation Rules** check box.

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

Note To configure dynamic PAT using a PAT pool instead of a single address, see the "Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool" section on page 40-4.

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

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

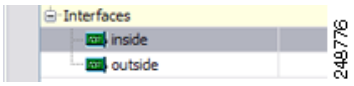

If you specify an interface name, then you enable *interface PAT*, where the specified interface IP address is used as the mapped address. To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** checkbox. With interface PAT, the NAT rule only applies to the specified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

Note You cannot specify an interface in transparent mode.

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

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

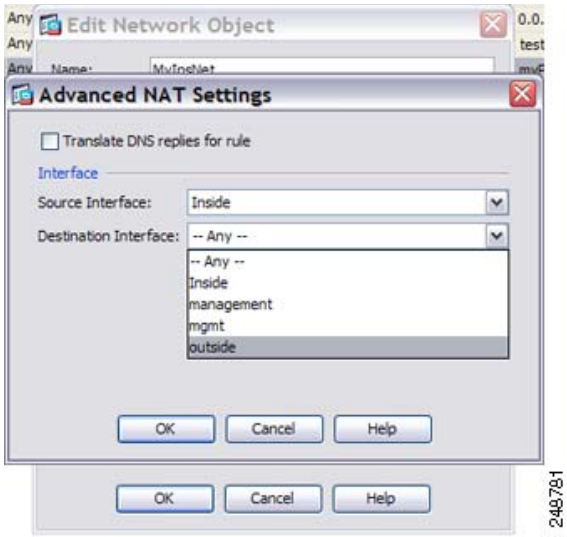

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See the "DNS and NAT" section on page 39-30 for more information.
- (Required for Transparent Firewall Mode) Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
- (Required for Transparent Firewall Mode) Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

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

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

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

This section describes how to configure a static NAT rule using network object NAT. For more information, see the "Static NAT" section on page 39-3.

Detailed Steps

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

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

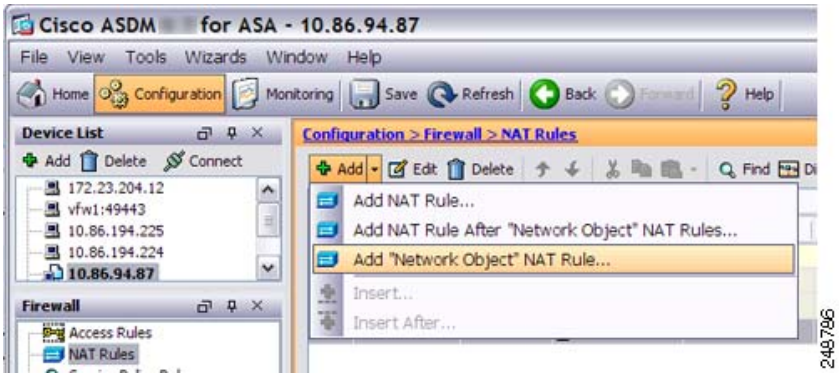

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the "Configuring a Network Object" section on page 25-3.

The Add/Edit Network Object dialog box appears.

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

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

b. Type—Network, Host, or Range.

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

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

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

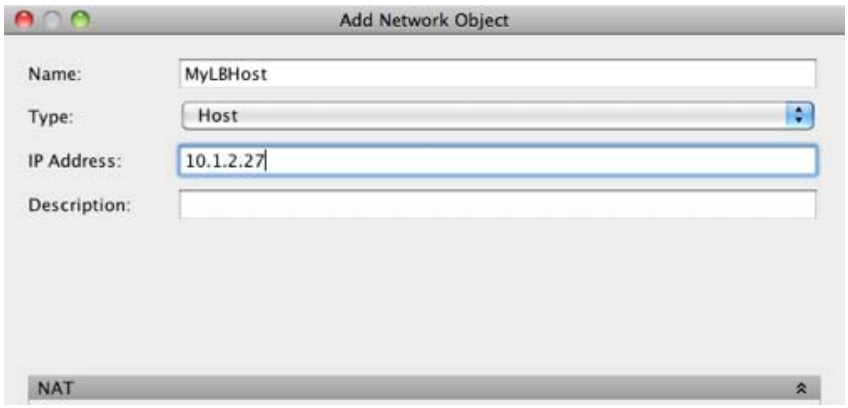

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

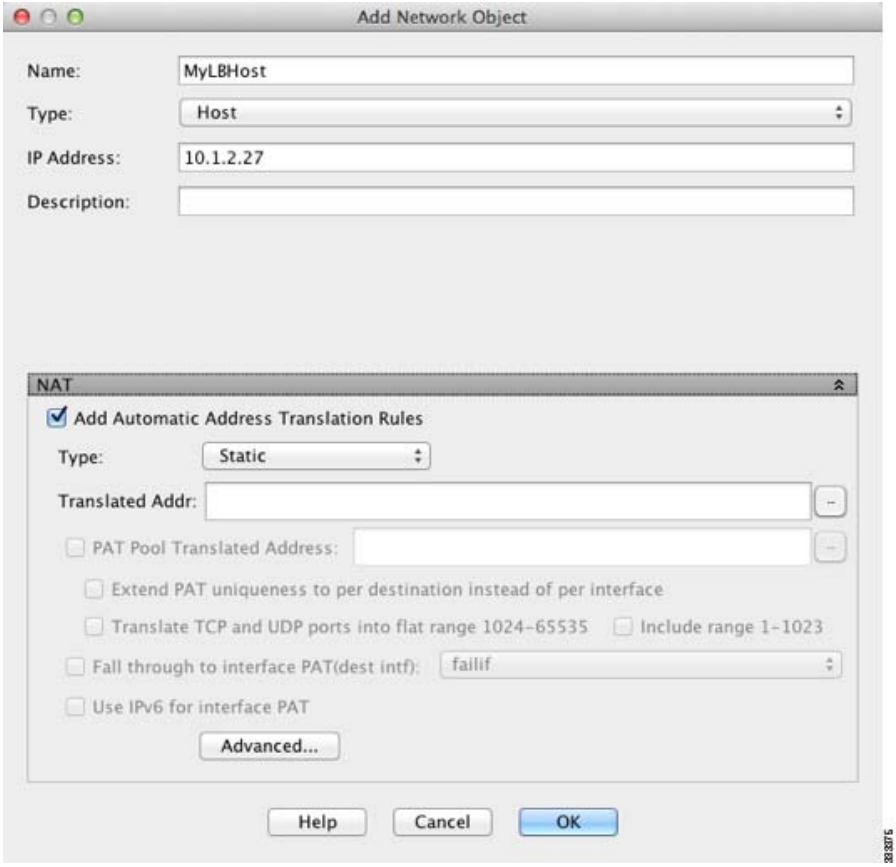

Step 4 Check the **Add Automatic Translation Rules** check box.

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

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

- Type an IP address.

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

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

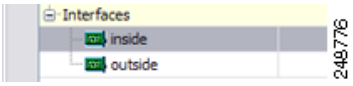

  To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** checkbox. Be sure to also configure a service on the Advanced NAT Settings dialog box (see Step 8). (You cannot specify an interface in transparent mode).

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

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

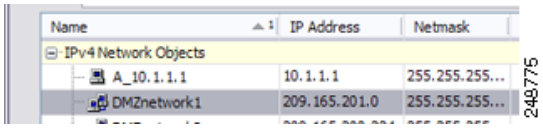

Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses. For more information, see the "Static NAT" section on page 39-3.

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

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

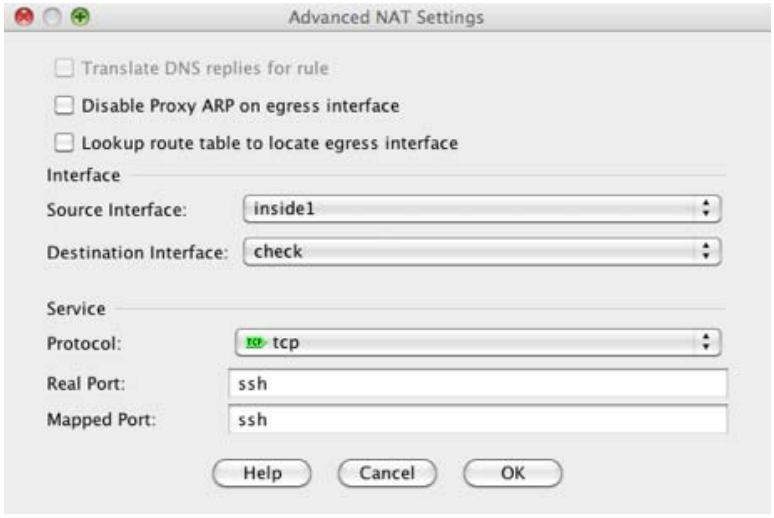

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See the "DNS and NAT" section on page 39-30 for more information.
- Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See the "Mapped Addresses and Routing" section on page 39-22 for more information.
- (Required for Transparent Firewall Mode) Interface:
  - Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
  - Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.
- Service:
  - Protocol—Configures static NAT-with-port-translation. Choose **tcp** or **udp**.
  - Real Port—You can type either a port number or a well-known port name (such as "ftp").
  - Mapped Port—You can type either a port number or a well-known port name (such as "ftp").

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

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

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

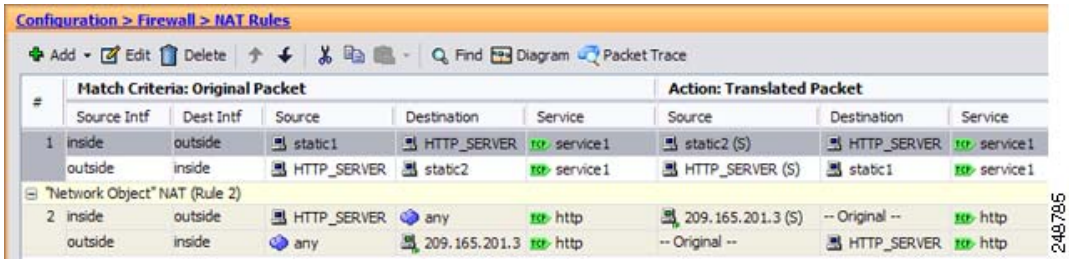

Configuring Identity NAT

This section describes how to configure an identity NAT rule using network object NAT. For more information, see the "Identity NAT" section on page 39-12.

Detailed Steps

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

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

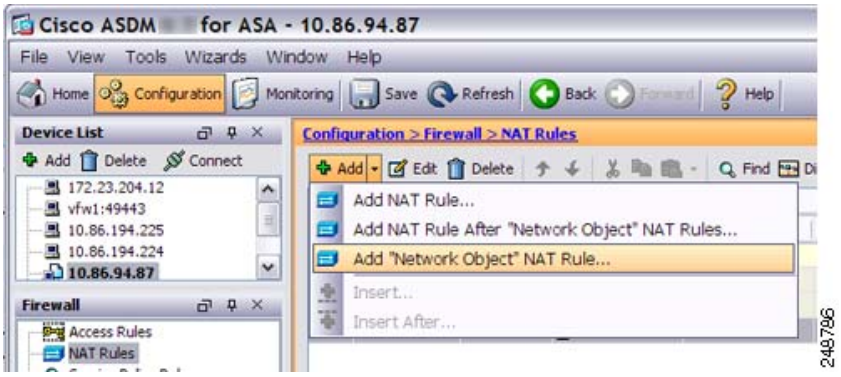

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the "Configuring a Network Object" section on page 25-3.

The Add/Edit Network Object dialog box appears.

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

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

b. Type—Network, Host, or Range.

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

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

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

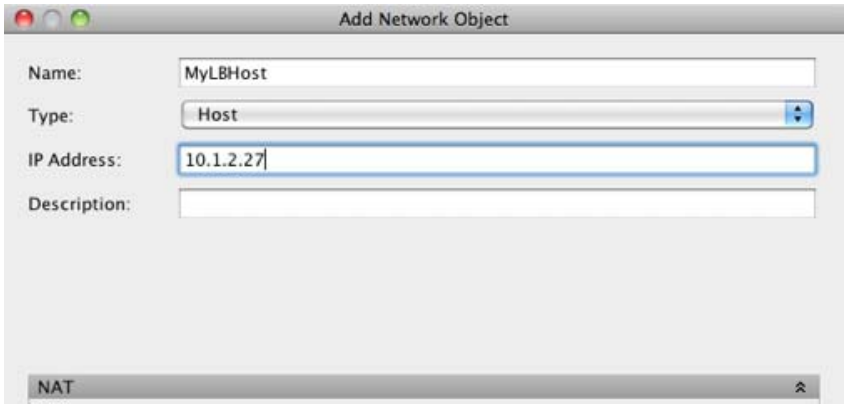

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

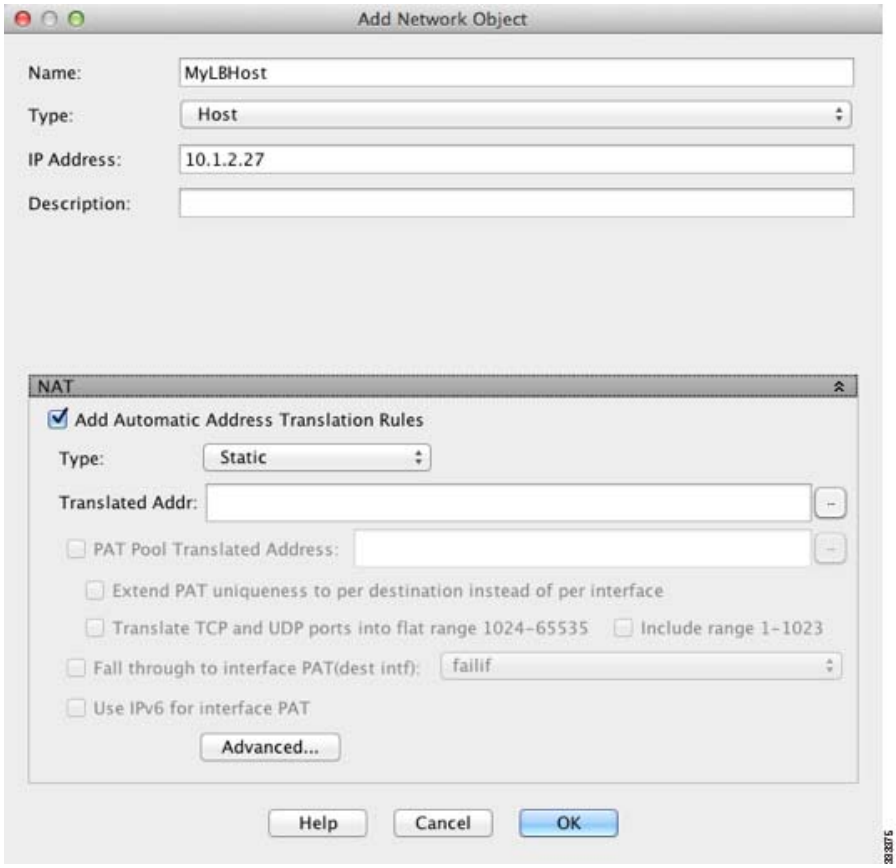

Step 4 Check the **Add Automatic Translation Rules** check box.

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

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

- Type the same IP address that you used for the real address.
- Click the browse button, and choose a network object with a matching IP address definition from the Browse Translated Addr dialog box.
- Click the browse button, and create a new network object with a matching IP address definition from the Browse Translated Addr dialog box.

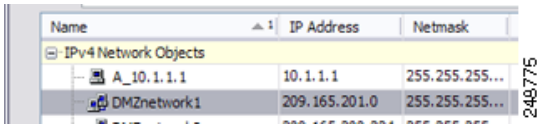

Step 7 (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

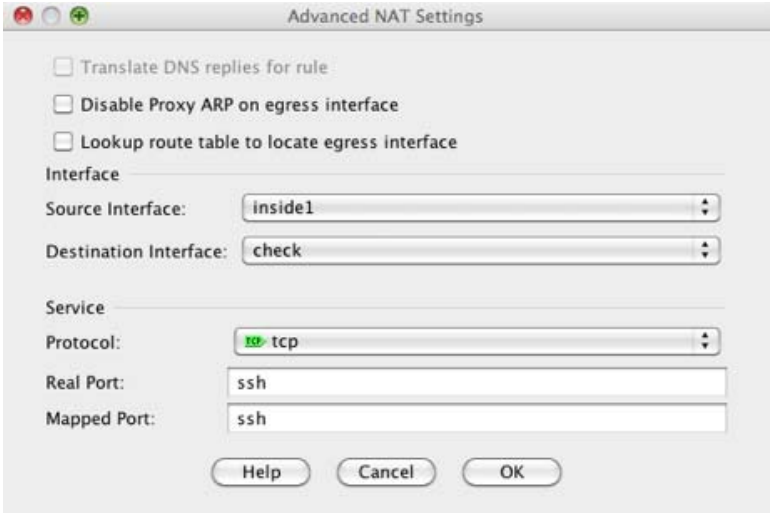

- Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See the "Mapped Addresses and Routing" section on page 39-22 for more information.
- (Routed mode; interface(s) specified) Lookup route table to locate egress interface—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See the "Determining the Egress Interface" section on page 39-24 for more information.
- (Required for Transparent Firewall Mode) Interface:
  - Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
  - Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

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

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

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

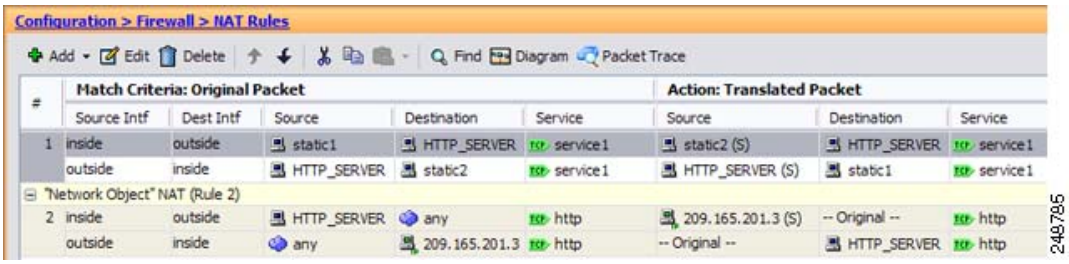

Configuring Per-Session PAT Rules

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT. For more information about per-session vs. multi-session PAT, see the "Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later)" section on page 39-11.

Defaults

By default, the following rules are installed:

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

These rules do not appear in the rule table.

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

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

Detailed Steps

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

Step 2 Click **Permit** or **Deny**.

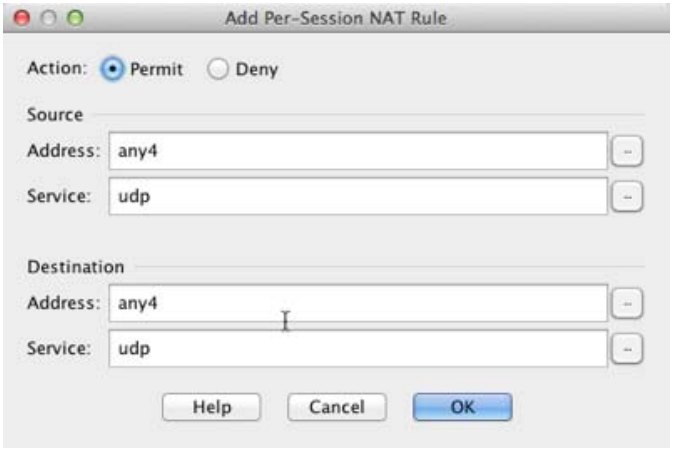

A permit rule uses per-session PAT; a deny rule uses multi-session PAT.

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

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

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

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

Step 7 Click **OK**.

Step 8 Click **Apply**.

Monitoring Network Object NAT

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

Fields

- Available Graphs—Lists the components you can graph.
  - Xlate Utilization—Displays the ASA NAT utilization.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected entry from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

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

Fields

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

Configuration Examples for Network Object NAT

This section includes the following configuration examples:

- Providing Access to an Inside Web Server (Static NAT), page 40-21
- NAT for Inside Hosts (Dynamic NAT) and NAT for an Outside Web Server (Static NAT), page 40-23
- Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many), page 40-28
- Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation), page 40-32
- DNS Server on Mapped Interface, Web Server on Real Interface (Static NAT with DNS Modification), page 40-35
- DNS Server and FTP Server on Mapped Interface, FTP Server is Translated (Static NAT with DNS Modification), page 40-38
- IPv4 DNS Server and FTP Server on Mapped Interface, IPv6 Host on Real Interface (Static NAT64 with DNS64 Modification), page 40-40

Providing Access to an Inside Web Server (Static NAT)

The following example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address. (See Figure 40-1).

Figure 40-1 Static NAT for an Inside Web Server

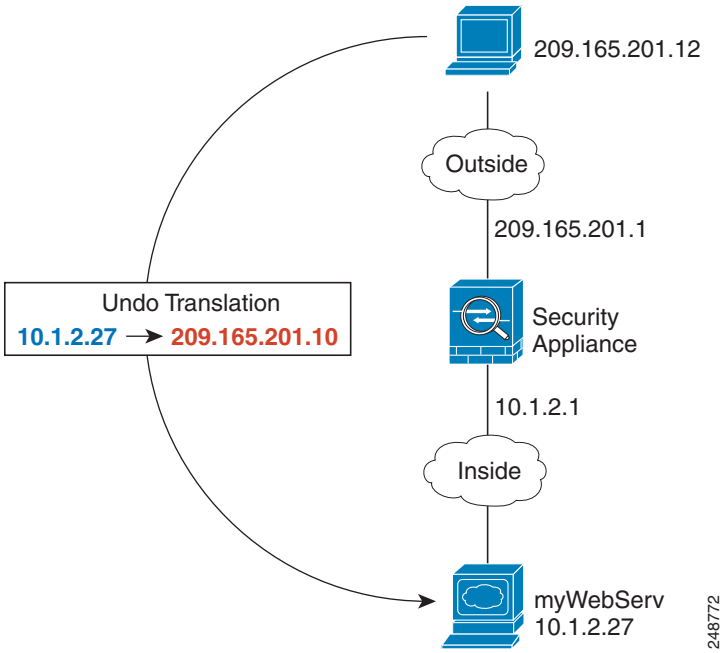

Step 1 Create a network object for the internal web server:

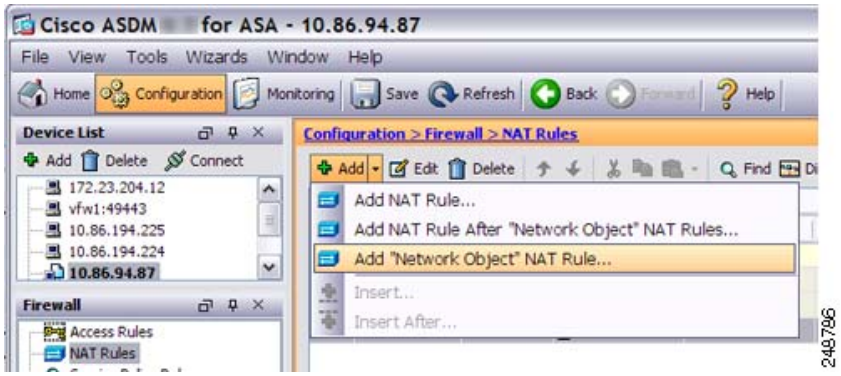

Step 2 Define the web server address:

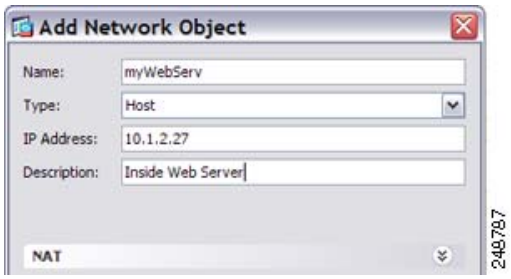

Step 3 Configure static NAT for the object:

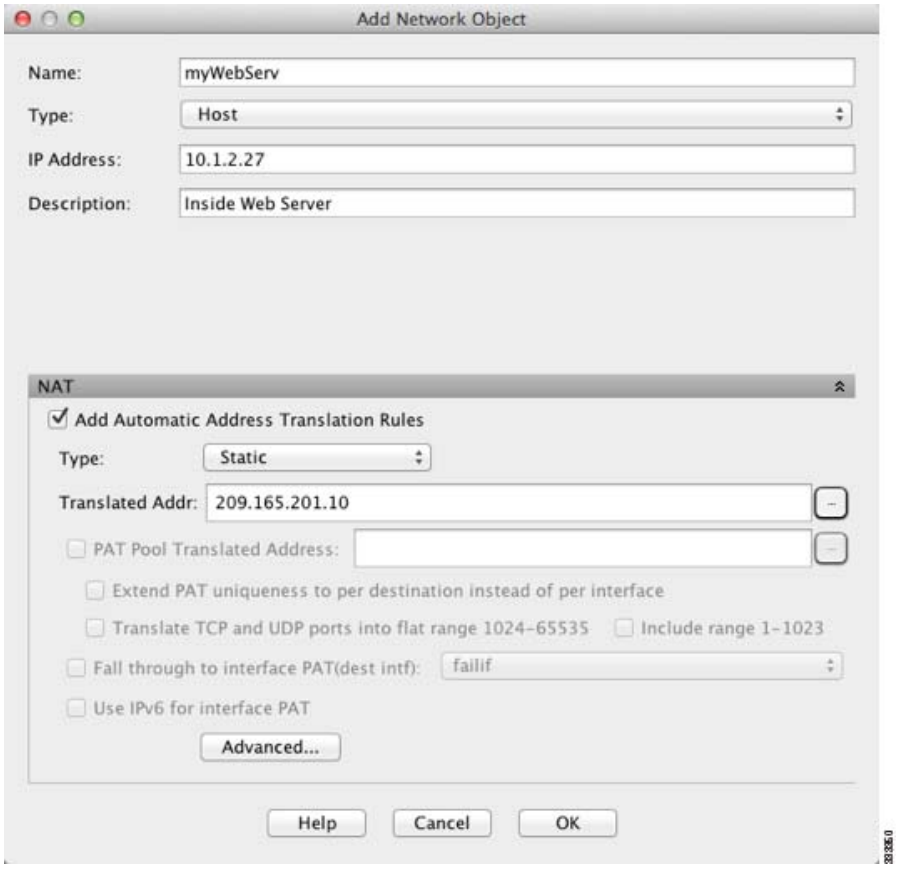

Step 4 Configure the real and mapped interfaces by clicking **Advanced**:

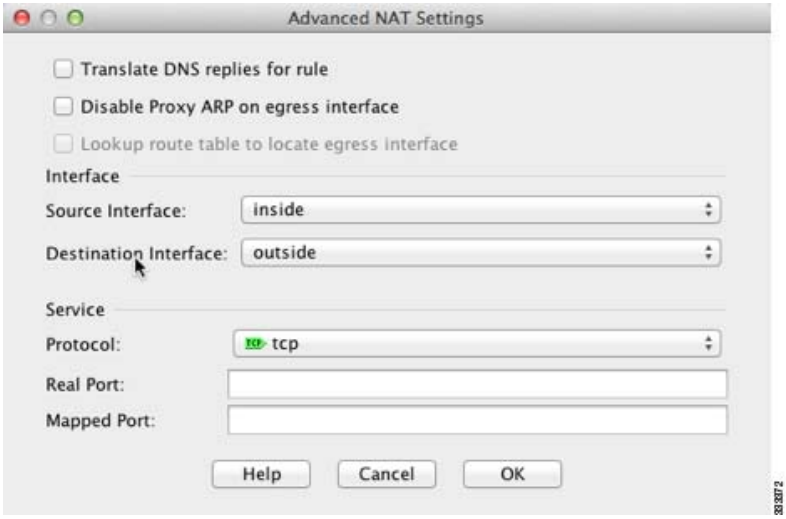

Step 5 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

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

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network. (See Figure 40-2).

Figure 40-2 Dynamic NAT for Inside, Static NAT for Outside Web Server

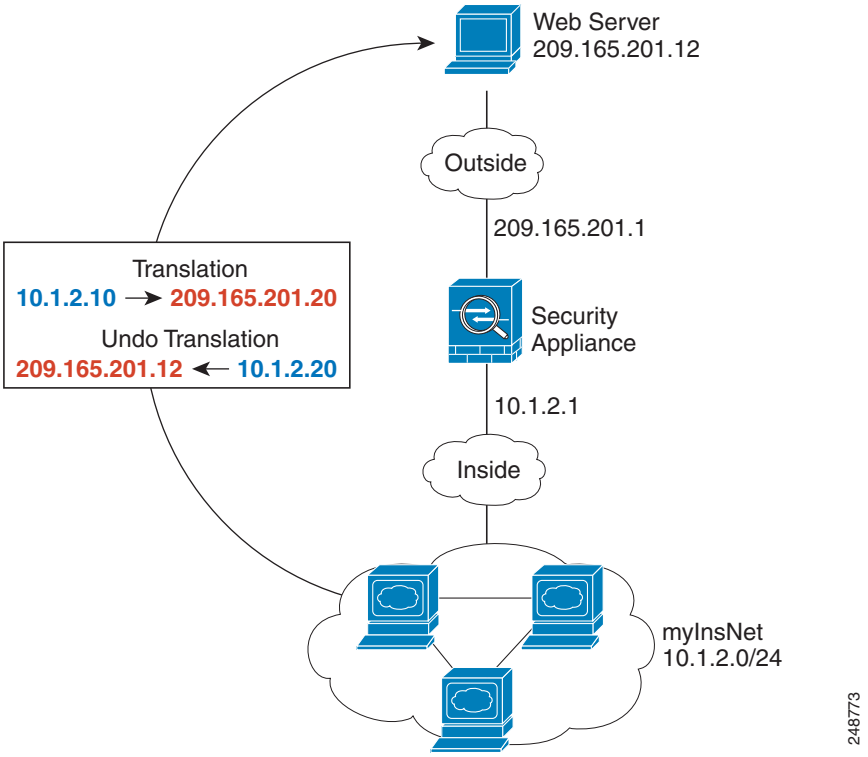

Step 1 Create a network object for the inside network:

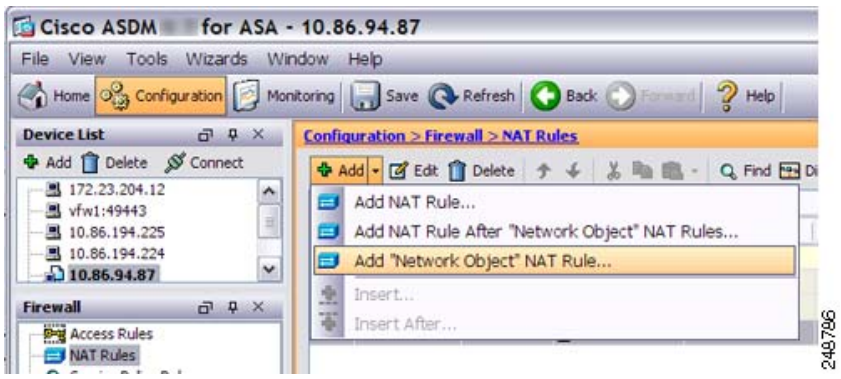

Step 2 Define the addresses for the inside network:

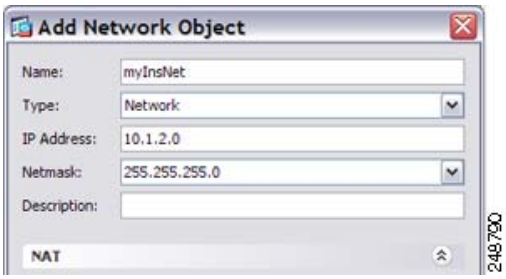

Step 3 Enable dynamic NAT for the inside network:

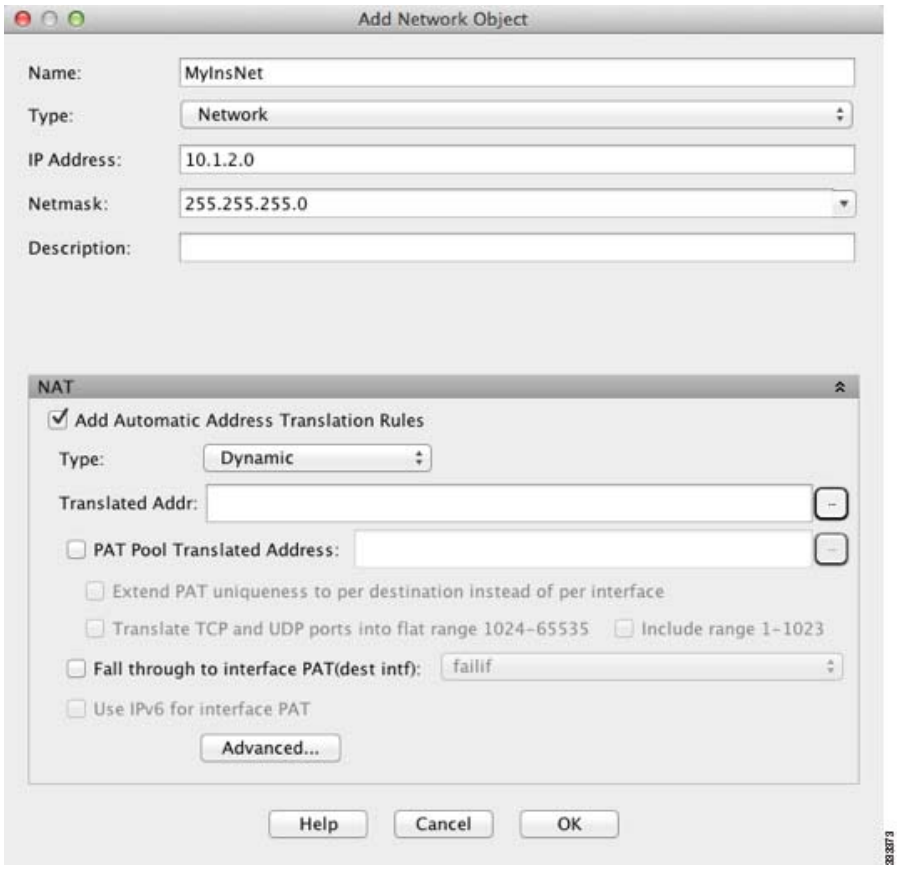

Step 4 For the Translated Addr field, add a new network object for the dynamic NAT pool to which you want to translate the inside addresses by clicking the browse button.

a. Add the new network object.

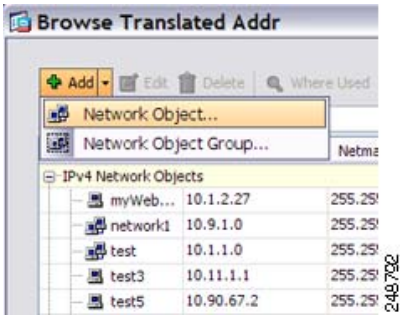

b. Define the NAT pool addresses, and click **OK**.

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

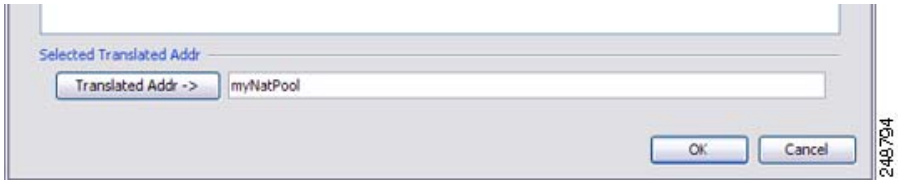

Step 5 Configure the real and mapped interfaces by clicking **Advanced**:

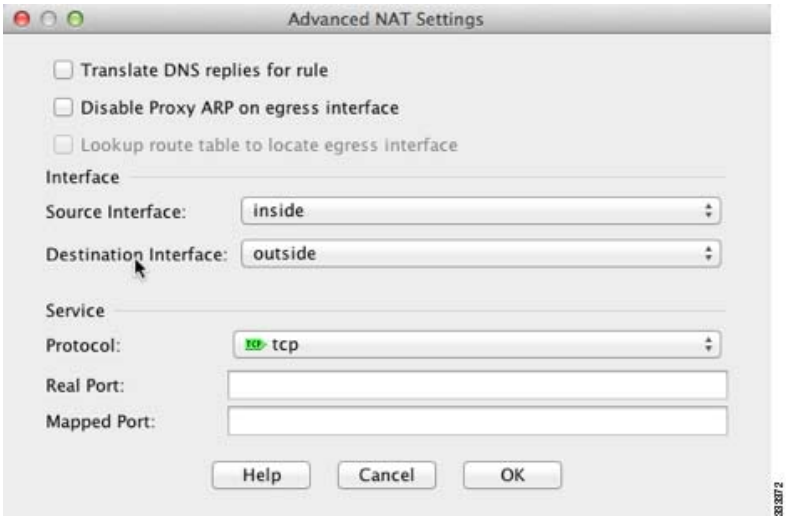

Step 6 Click **OK** to return to the Edit Network Object dialog box, click then click **OK** again to return to the NAT Rules table.

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

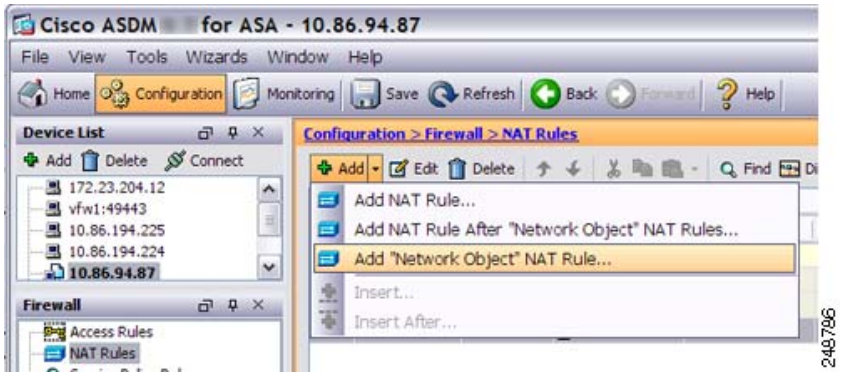

Step 8 Define the web server address:

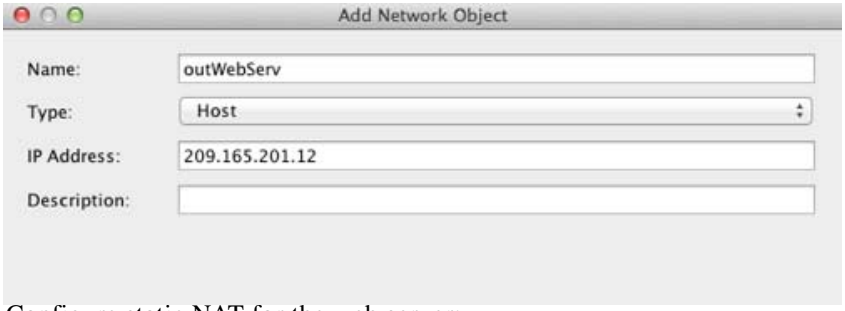

Step 9 Configure static NAT for the web server:

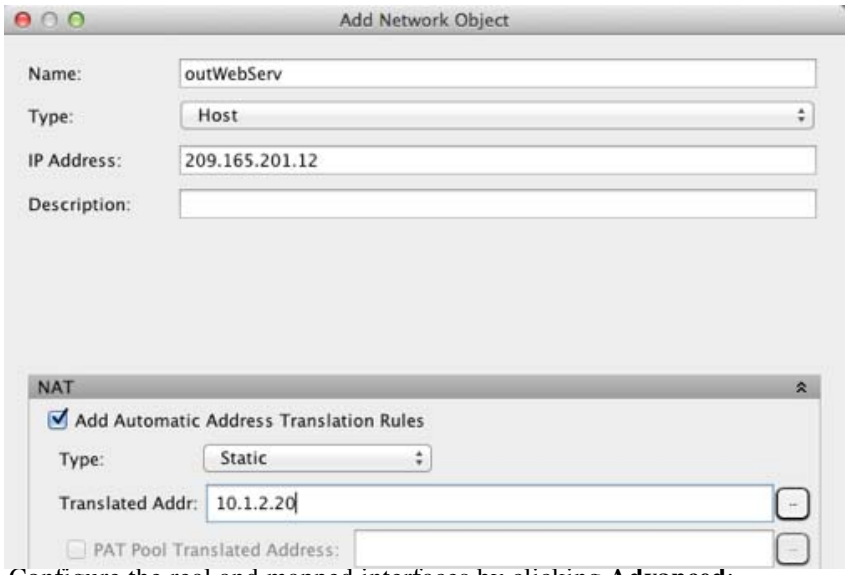

Step 10 Configure the real and mapped interfaces by clicking **Advanced**:

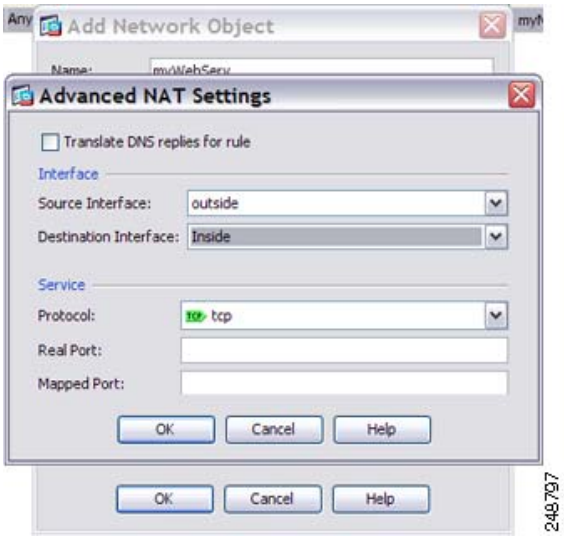

Step 11 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

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

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server. (See Figure 40-3).

Figure 40-3 Static NAT with One-to-Many for an Inside Load Balancer

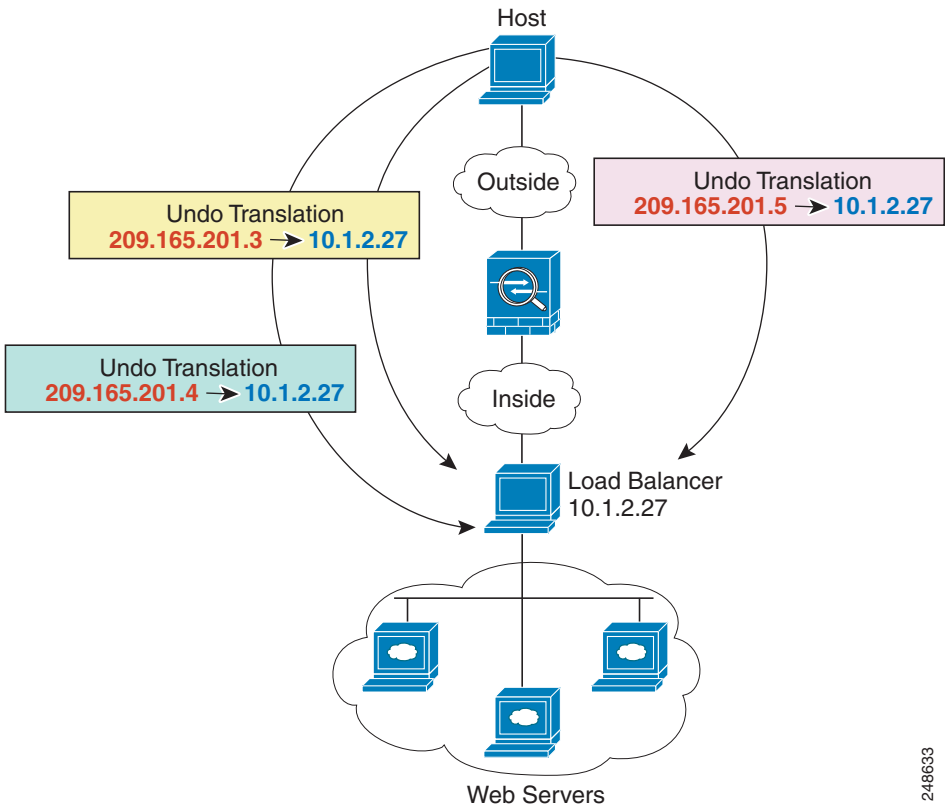

Step 1 Create a network object for the load balancer:

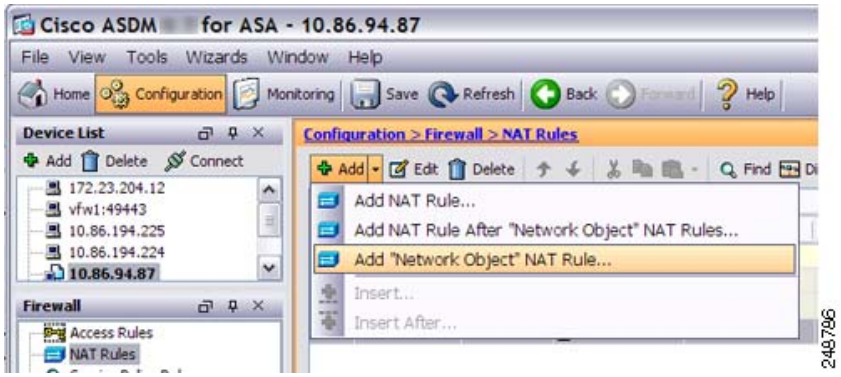

Step 2 Define the load balancer address:

Step 3 Configure static NAT for the load balancer:

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

a. Add the new network object.

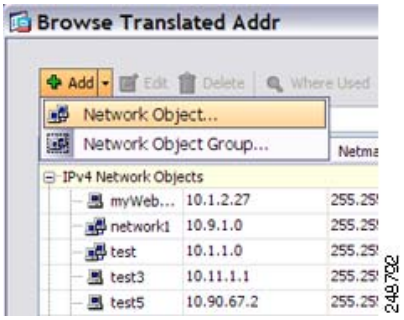

b. Define the static NAT group of addresses, and click **OK**.

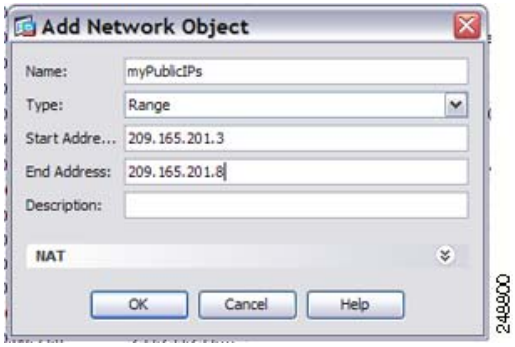

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5 Configure the real and mapped interfaces by clicking **Advanced**:

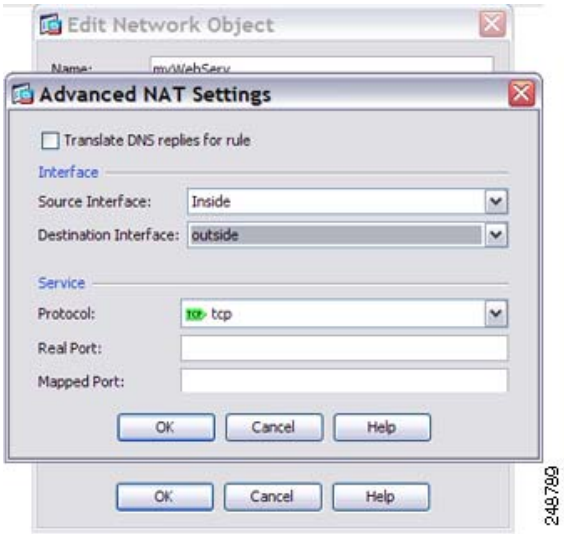

Step 6 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports. (See Figure 40-4.)

Figure 40-4 Static NAT-with-Port-Translation

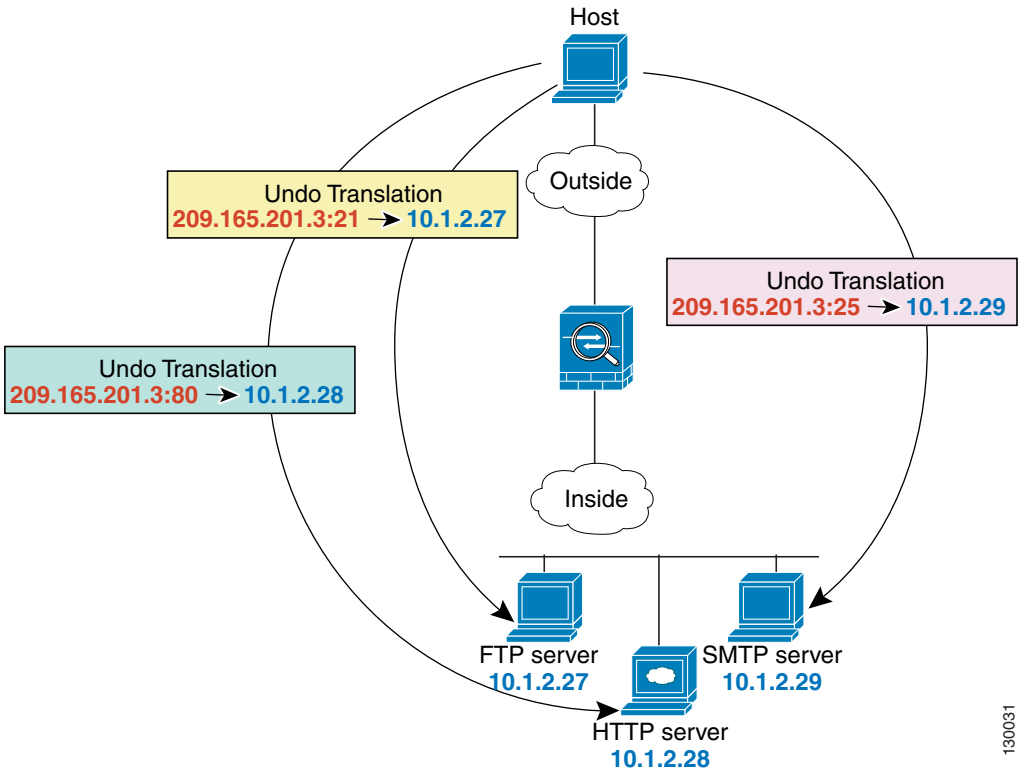

Step 1 Create a network object for the FTP server address:

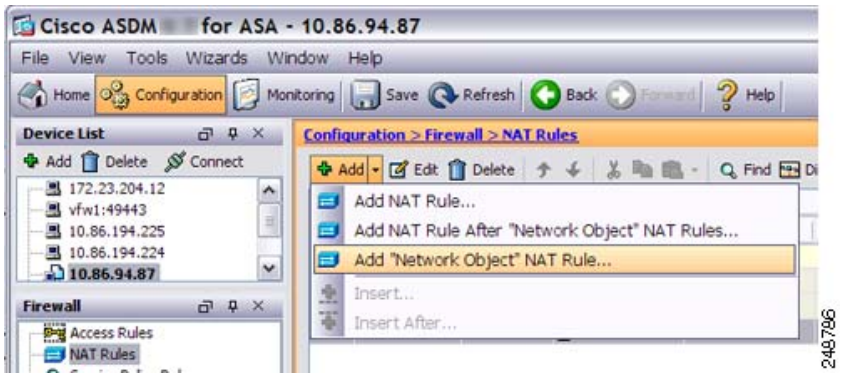

Step 2 Define the FTP server address, and configure static NAT with identity port translation for the FTP server:

Step 3 Click **Advanced** to configure the real and mapped interfaces and port translation for FTP.

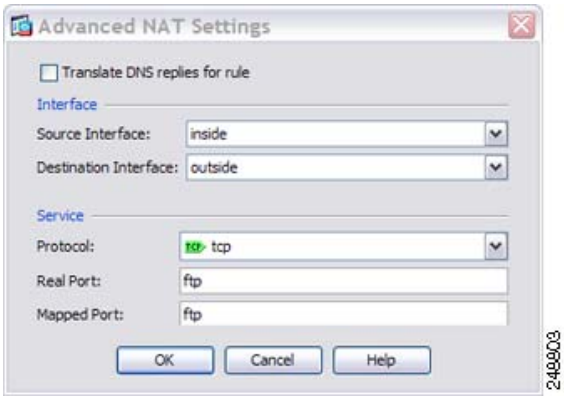

Step 4 Create a network object for the HTTP server address:

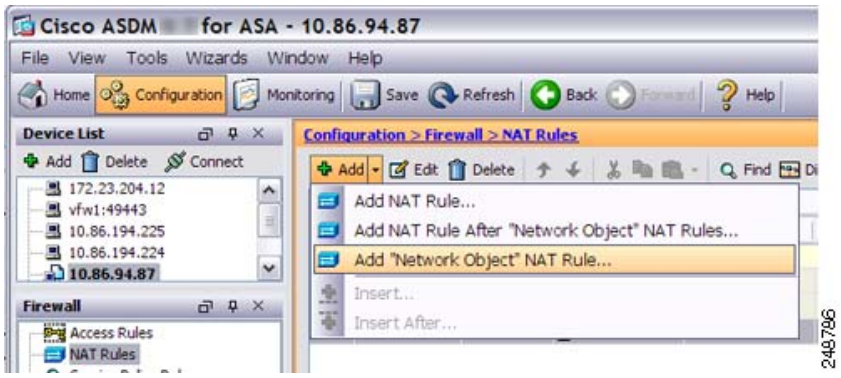

Step 5 Define the HTTP server address, and configure static NAT with identity port translation for the HTTP server:

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

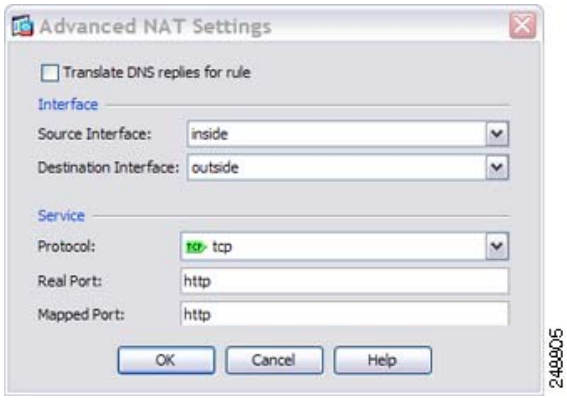

Step 7 Create a network object for the SMTP server address:

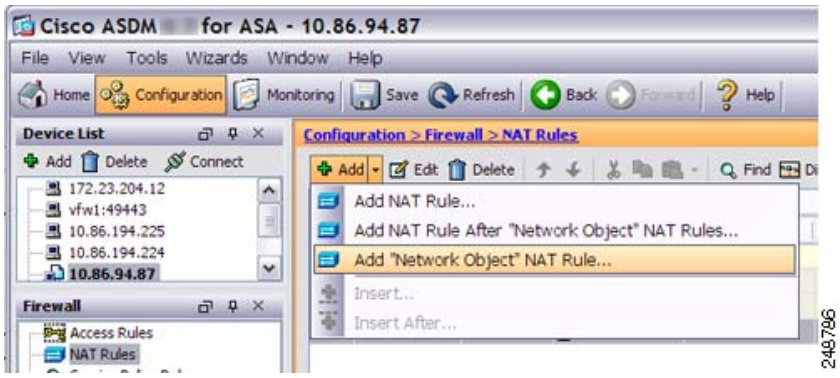

Step 8 Define the SMTP server address, and configure static NAT with identity port translation for the SMTP server:

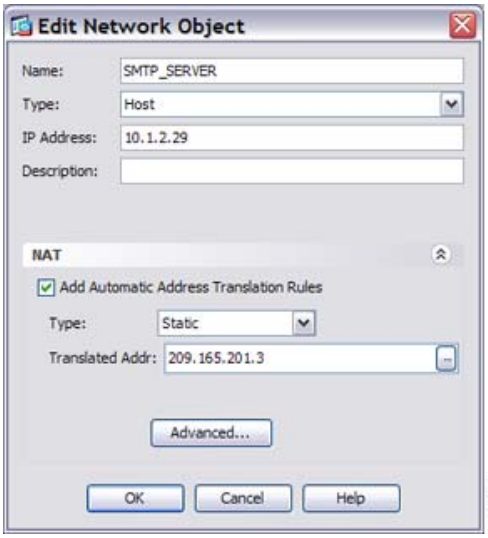

Step 9 Click **Advanced** to configure the real and mapped interfaces and port translation for SMTP.

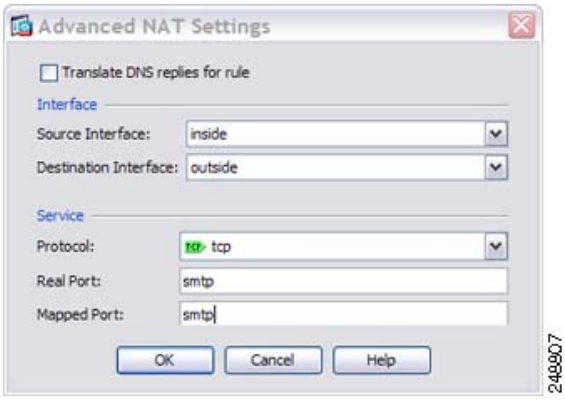

Step 10 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

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

For example, a DNS server is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network. (See Figure 40-5.) In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The ASA refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

Figure 40-5 DNS Reply Modification

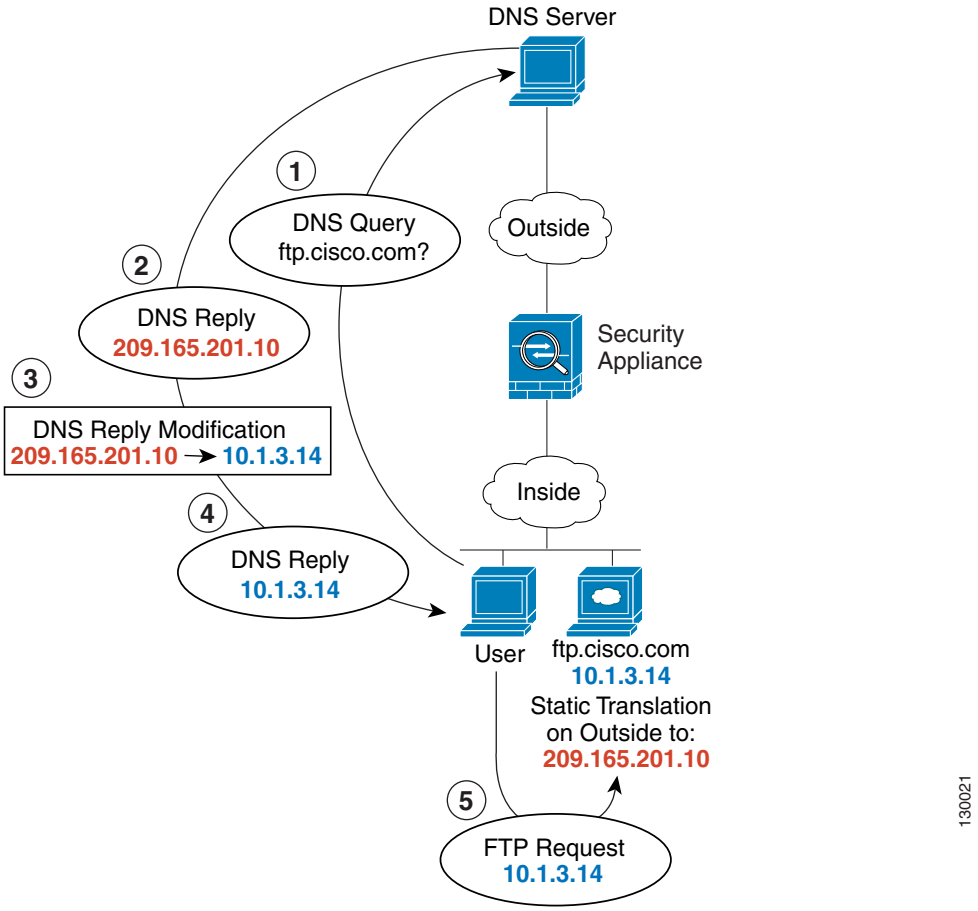

Step 1 Create a network object for the FTP server address:

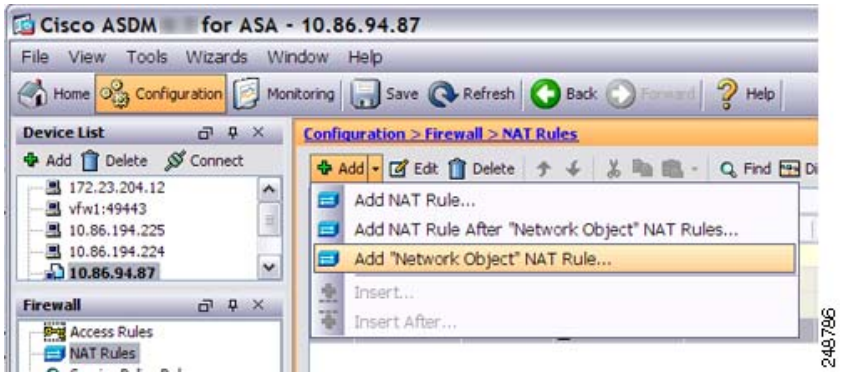

Step 2 Define the FTP server address, and configure static NAT with DNS modification:

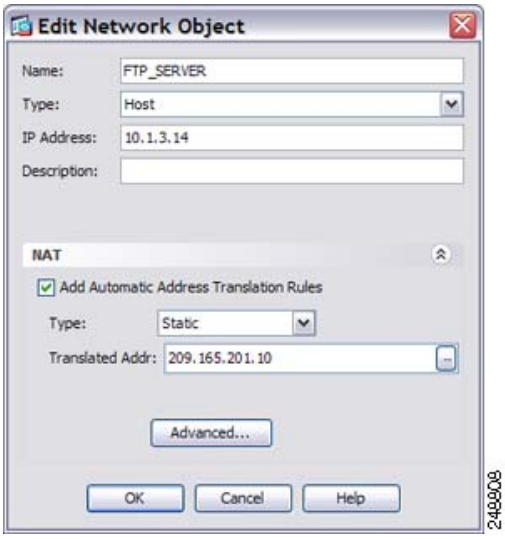

Step 3 Click **Advanced** to configure the real and mapped interfaces and DNS modification.

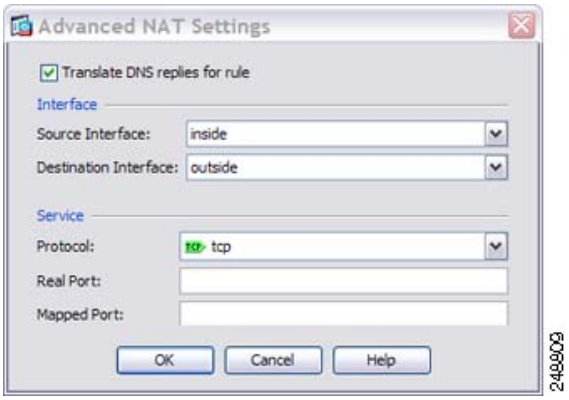

Step 4 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

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

Figure 40-6 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.201.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 40-6 DNS Reply Modification Using Outside NAT

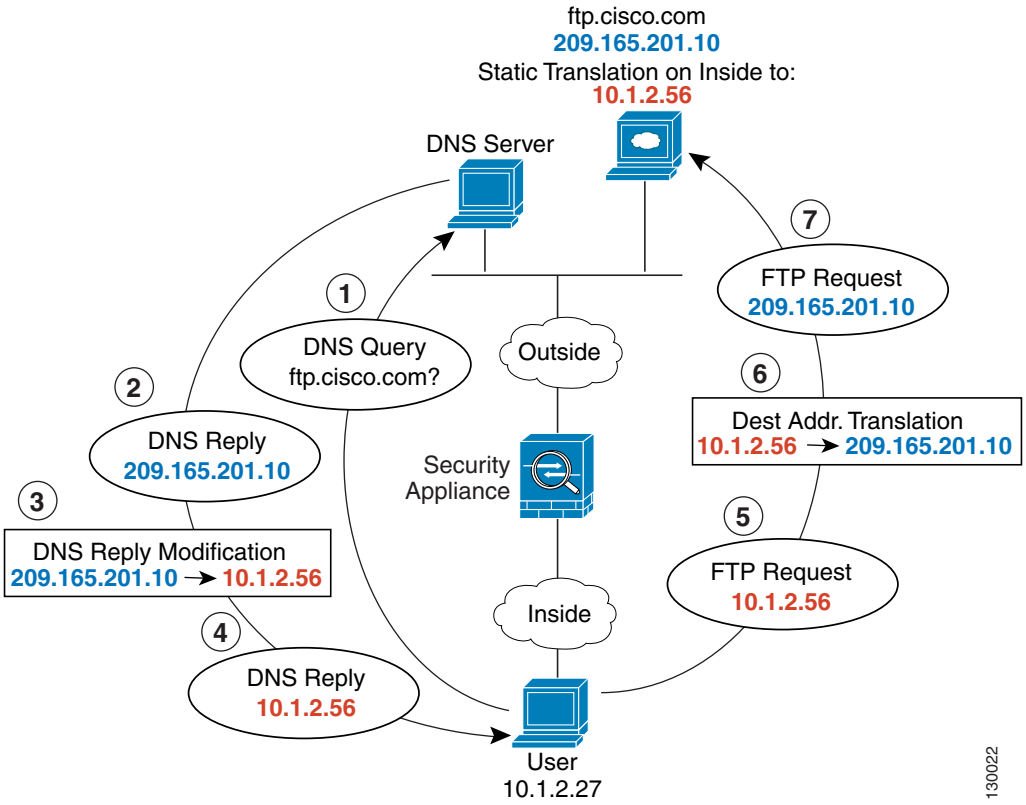

Step 1 Create a network object for the FTP server address:

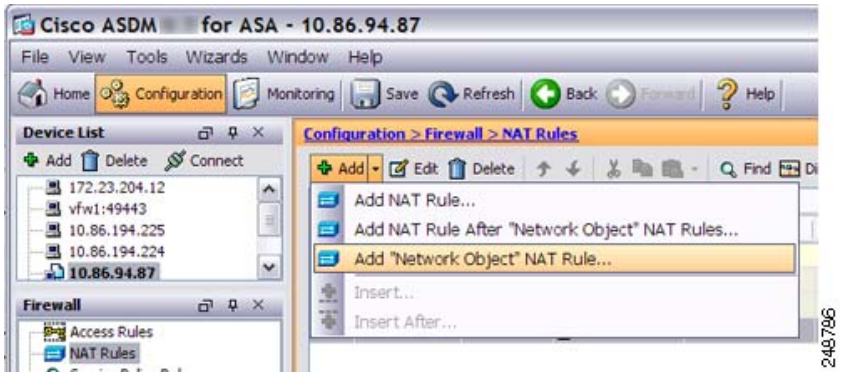

Step 2 Define the FTP server address, and configure static NAT with DNS modification:

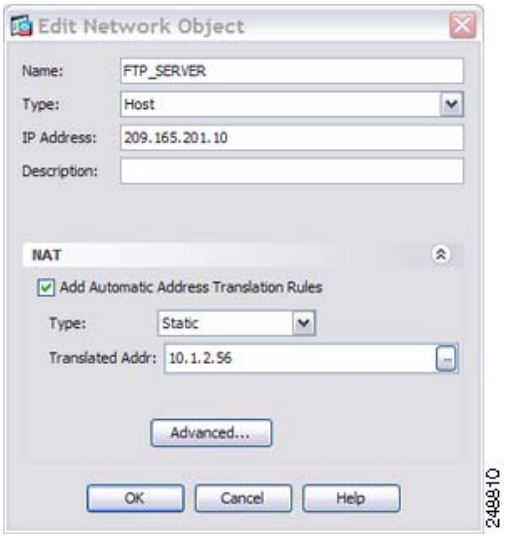

Step 3 Click **Advanced** to configure the real and mapped interfaces and DNS modification.

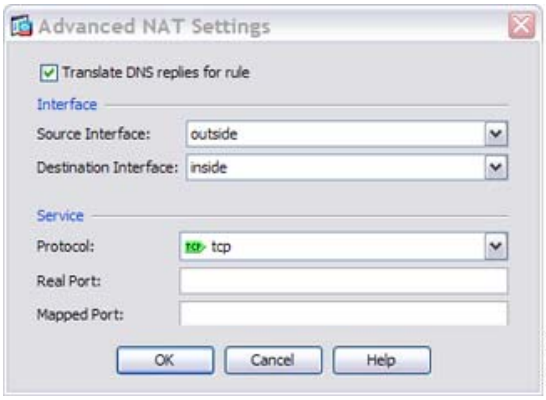

Step 4 Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

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

Figure 40-6 shows an FTP server and DNS server on the outside IPv4 network. The ASA has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225. Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.

Figure 40-7 DNS Reply Modification Using Outside NAT

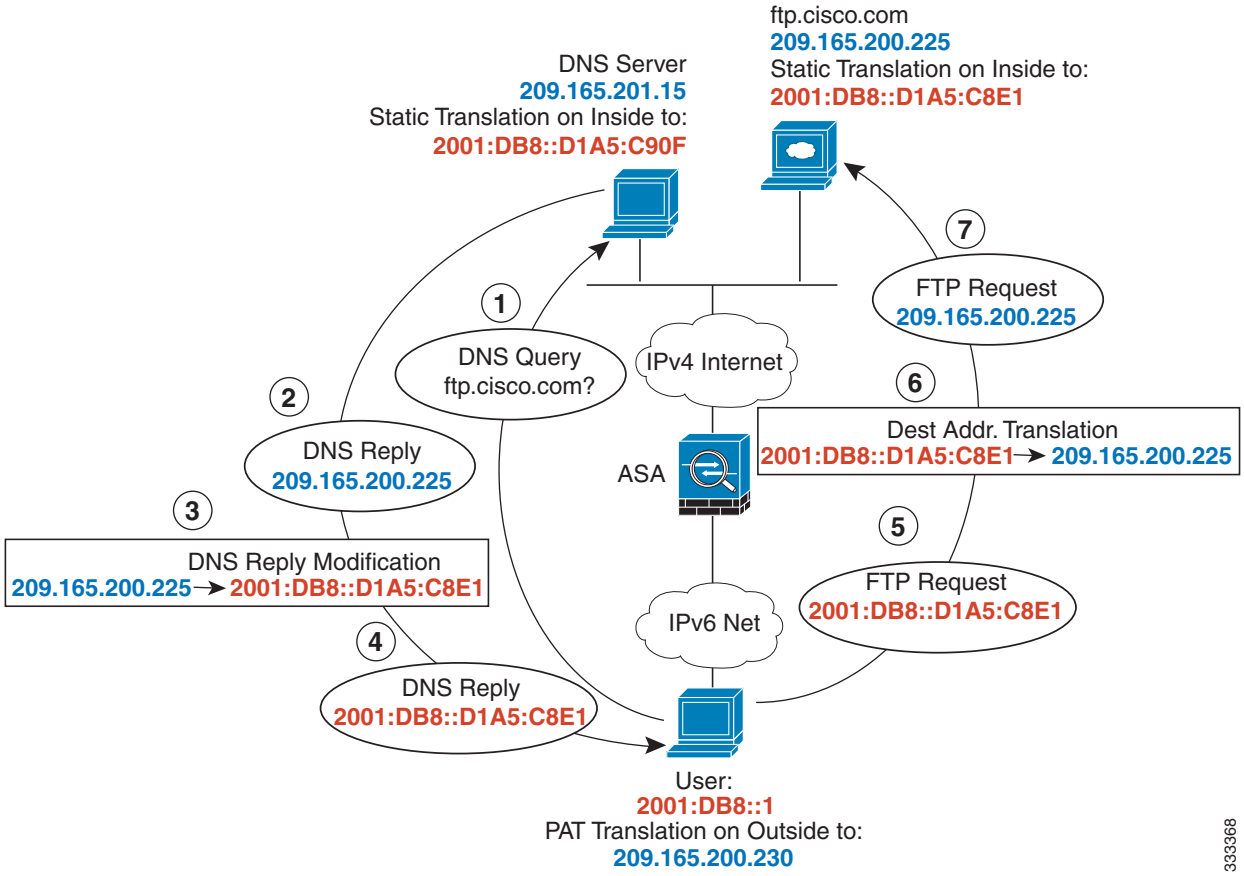

Step 1 Configure static NAT with DNS modification for the FTP server.

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

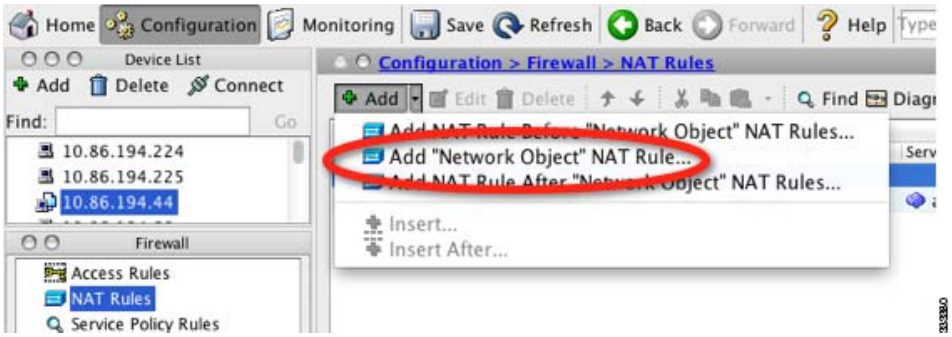

b. Define the FTP server address, and configure static NAT with DNS modification and, because this is a one-to-one translation, configure the one-to-one method for NAT46.

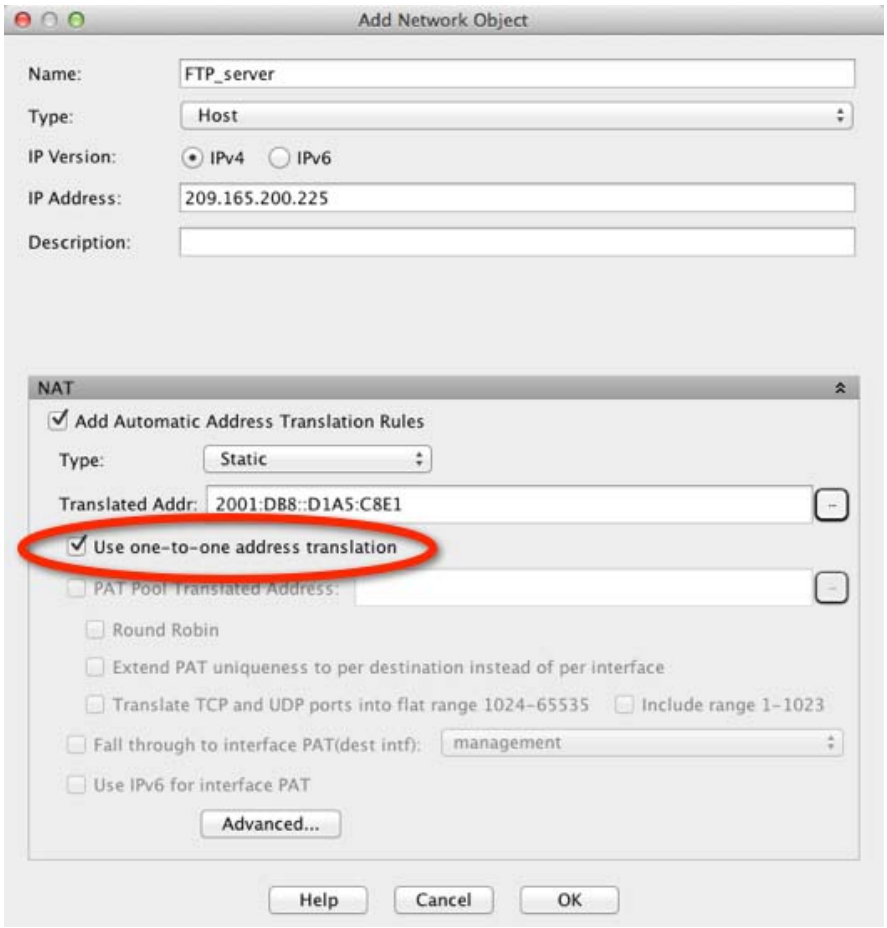

c. Click **Advanced** to configure the real and mapped interfaces and DNS modification.

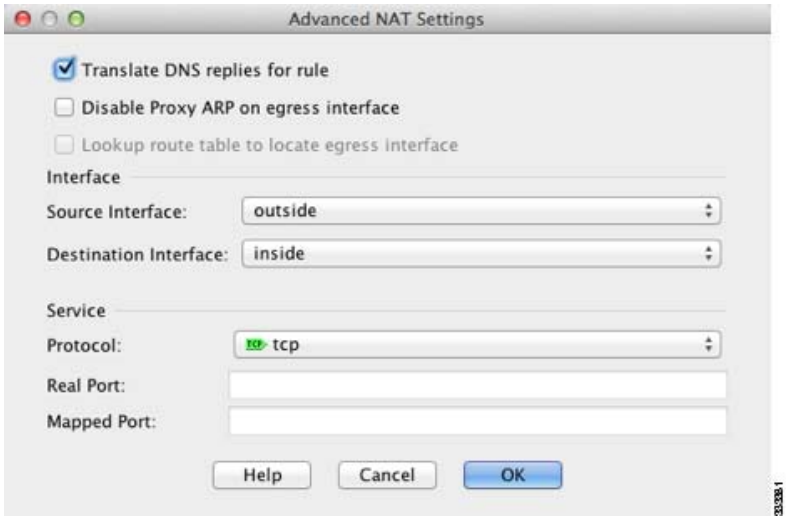

d. Click **OK** to return to the Edit Network Object dialog box.

Step 2 Configure NAT for the DNS server.

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

b. Define the DNS server address, and configure static NAT using the one-to-one method.

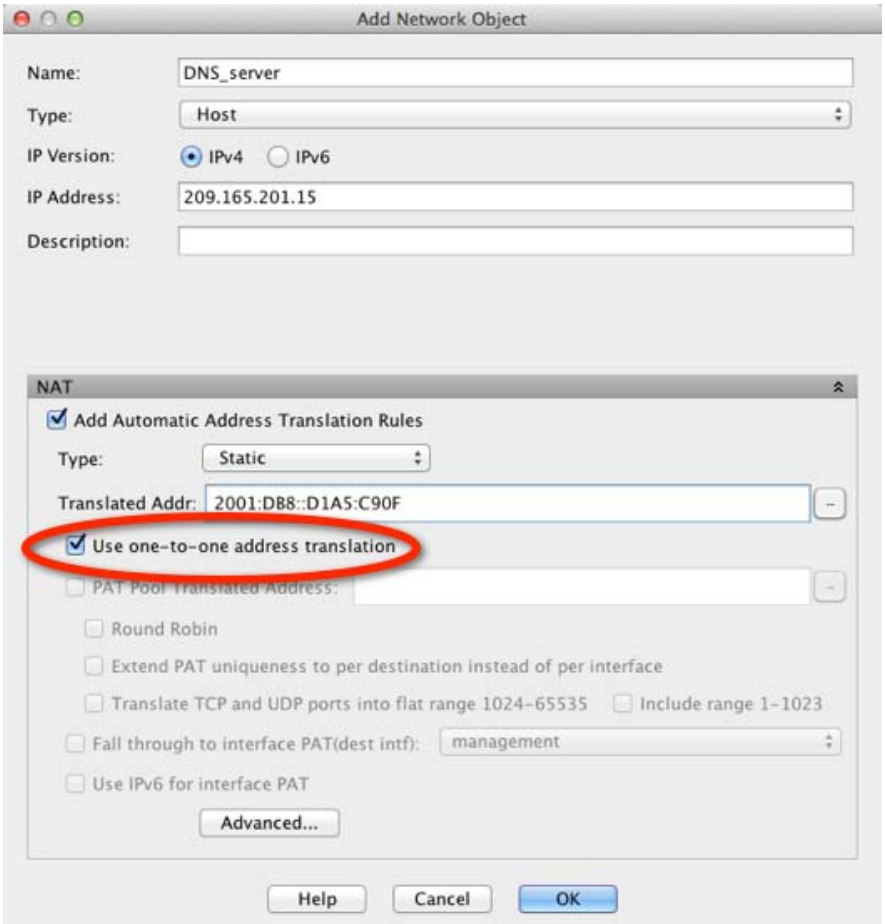

c. Click **Advanced** to configure the real and mapped interfaces.

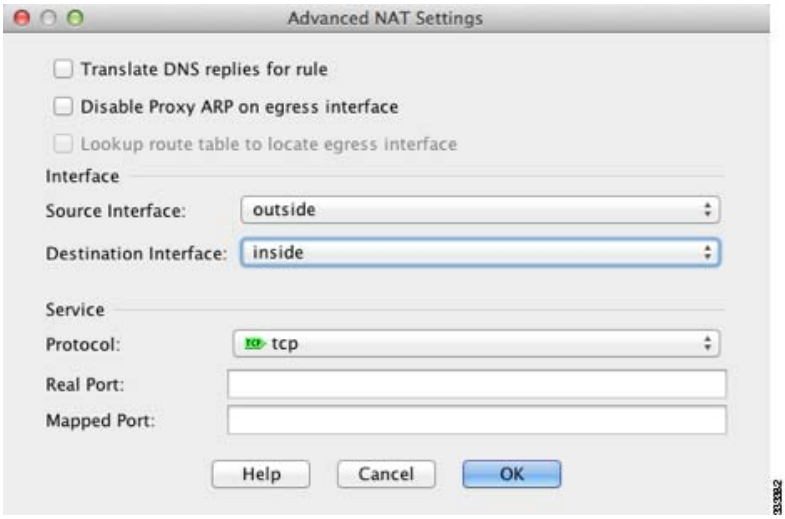

d. Click **OK** to return to the Edit Network Object dialog box.

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

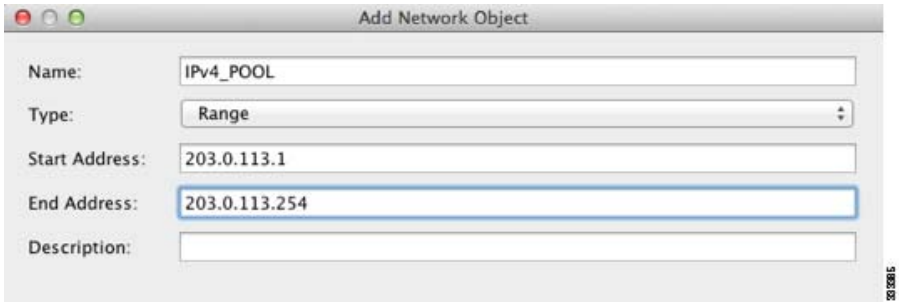

Under NAT, uncheck the **Add Automatic Address Translation Rules** check box.

Step 4 Configure PAT for the inside IPv6 network.

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

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

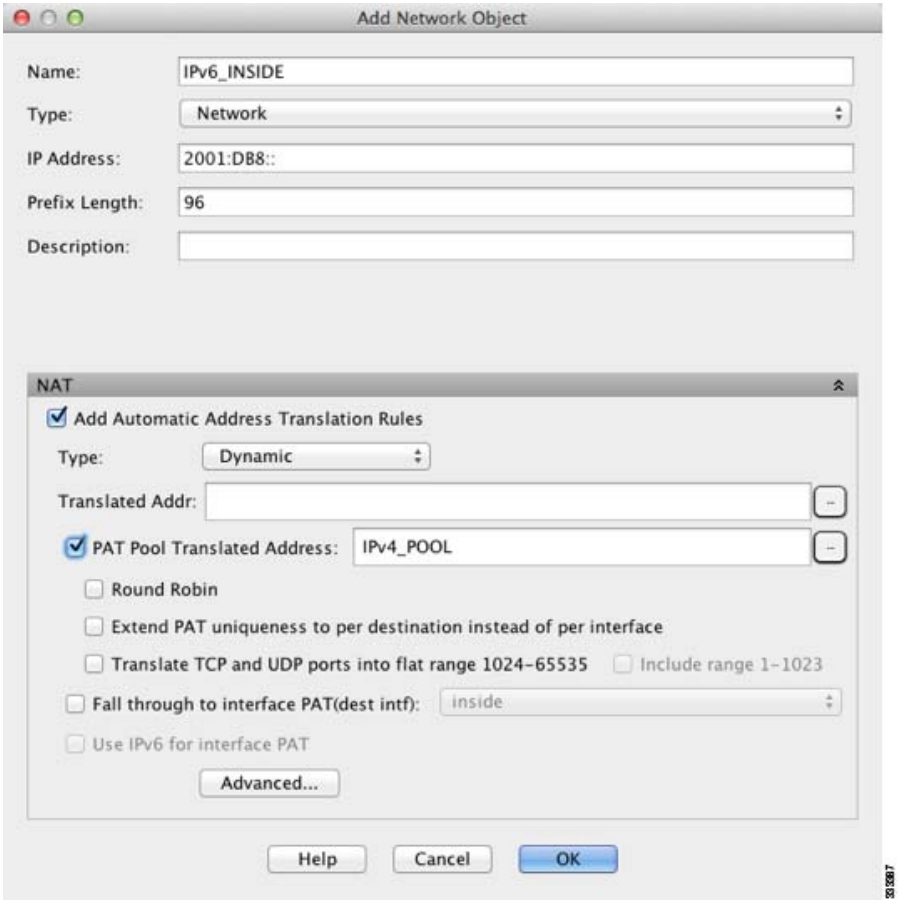

c. Next to the PAT Pool Translated Address field, click the **...** button to choose the PAT pool you created earlier, and click **OK**.

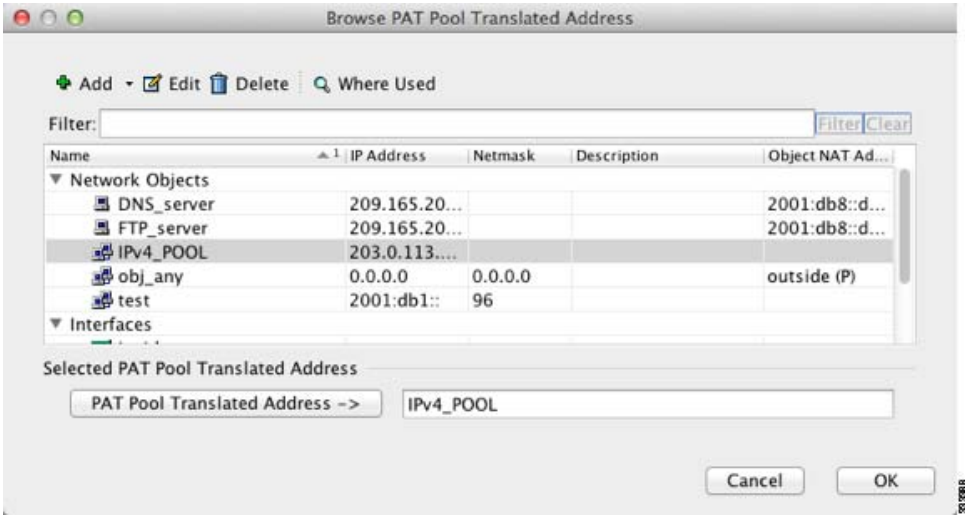

d. Click **Advanced** to configure the real and mapped interfaces.

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

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

Feature History for Network Object NAT

Table 40-1 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 40-1 Feature History for Network Object NAT

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

Table 40-1 Feature History for Network Object NAT (continued)

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

Table 40-1 Feature History for Network Object NAT (continued)

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

Table 40-1 Feature History for Network Object NAT (continued)

Feature Name	Platform Releases	Feature Information
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode. We modified the following screen: Configuration > Firewall > Objects > Network Objects/Group.
Per-session PAT	9.0(1)	The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.

Configuring Twice NAT (ASA 8.3 and Later)

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

- Information About Twice NAT, page 41-1
- Licensing Requirements for Twice NAT, page 41-2
- Prerequisites for Twice NAT, page 41-2
- Guidelines and Limitations, page 41-2
- Default Settings, page 41-4
- Configuring Twice NAT, page 41-4
- Monitoring Twice NAT, page 41-30
- Configuration Examples for Twice NAT, page 41-30
- Feature History for Twice NAT, page 41-48

Note For detailed information about how NAT works, see Chapter 39, "Information About NAT (ASA 8.3 and Later)."

Information About Twice NAT

Twice NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that a source address should be translated to A when going to destination X, but be translated to B when going to destination Y, for example.

Note For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then in the command, you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

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

For detailed information about the differences between twice NAT and network object NAT, see the "How NAT is Implemented" section on page 39-15.

Twice NAT rules are added to section 1 of the NAT rules table, or if specified, section 3. For more information about NAT ordering, see the "NAT Rule Order" section on page 39-20.

Licensing Requirements for Twice NAT

Model	License Requirement
All models	Base License.

Prerequisites for Twice NAT

- For both the real and mapped addresses, configure network objects or network object groups. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets. To create a network object or group, see the "Configuring Network Objects and Groups" section on page 25-2.
- For static NAT-with-port-translation, configure TCP or UDP service objects. To create a service object, see the "Configuring Service Objects and Service Groups" section on page 25-4.

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

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

- Supported in routed and transparent firewall mode.
- In transparent mode, you must specify the real and mapped interfaces; you cannot use --Any--.
- In transparent mode, you cannot configure interface PAT, because the transparent mode interfaces do not have IP addresses. You also cannot use the management IP address as a mapped address.
- In transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

IPv6 Guidelines

- Supports IPv6.

- For routed mode, you can also translate between IPv4 and IPv6.
- For transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.
- For transparent mode, a PAT pool is not supported for IPv6.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

Additional Guidelines

- You cannot configure FTP destination port translation when the source IP address is a subnet (or any other application that uses a secondary connection); the FTP data channel establishment does not succeed.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT information is used, you can clear the translation table using the **clear xlate** command. However, clearing the translation table disconnects all current connections that use translations.

Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** command. This safeguard ensures that the same address is not assigned to multiple hosts.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- When using the **any** keyword in a NAT rule, the definition of "any" traffic (IPv4 vs. IPv6) depends on the rule. Before the ASA performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the ASA can determine the value of **any** in a NAT rule. For example, if you configure a rule from "any" to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means "any IPv6 traffic." If you configure a rule from "any" to "any," and you map the source to the interface IPv4 address, then **any** means "any IPv4 traffic" because the mapped interface address implies that the destination is also IPv4.
- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.
- You can use the same objects in multiple rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify --Any-- interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), use the interface name instead of the IP address.
  - (Transparent mode) The management IP address.
  - (Dynamic NAT) The standby interface IP address when VPN is enabled.
  - Existing VPN pool addresses.

Default Settings

- By default, the rule is added to the end of section 1 of the NAT table.
- (Routed mode) The default real and mapped interface is Any, which applies the rule to all interfaces.
- (8.3(1), 8.3(2), and 8.4(1)) The default behavior for identity NAT has proxy ARP disabled. You cannot configure this setting. (8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired.
- If you specify an optional interface, then the ASA uses the NAT configuration to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration, but you have the option to always use a route lookup instead.

Configuring Twice NAT

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

- Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool, page 41-4
- Configuring Dynamic PAT (Hide), page 41-12
- Configuring Static NAT or Static NAT-with-Port-Translation, page 41-18
- Configuring Identity NAT, page 41-24
- Configuring Per-Session PAT Rules, page 41-29

Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool

This section describes how to configure twice NAT for dynamic NAT or for dynamic PAT using a PAT pool. For more information, see the "Dynamic NAT" section on page 39-8 or the "Dynamic PAT" section on page 39-10.

Guidelines

For a PAT pool:

- If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you have a lot of traffic that uses the lower port ranges, you can now specify for a PAT pool a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535.
- (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT and a flat range, then the other rule must also specify extended PAT and a flat range.

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

- Many application inspections do not support extended PAT. See the "Default Settings" section on page 57-4 in Chapter 57, "Getting Started with Application Layer Protocol Inspection," for a complete list of unsupported inspections.

- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.
- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.

For round robin for a PAT pool:

- (8.4(3) and later, not including 8.5(1) or 8.6(1)) If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. **Note**: This "stickiness" does not survive a failover. If the ASA fails over, then subsequent connections from a host may not use the initial IP address.
- (8.4(2), 8.5(1), and 8.6(1)) If a host has an existing connection, then subsequent connections from that host will likely use *different* PAT addresses for each connection because of the round robin allocation. In this case, you may have problems when accessing two websites that exchange information about the host, for example an e-commerce site and a payment site. When these sites see two different IP addresses for what is supposed to be a single host, the transaction may fail.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Detailed Steps

To configure dynamic NAT, perform the following steps:

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

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

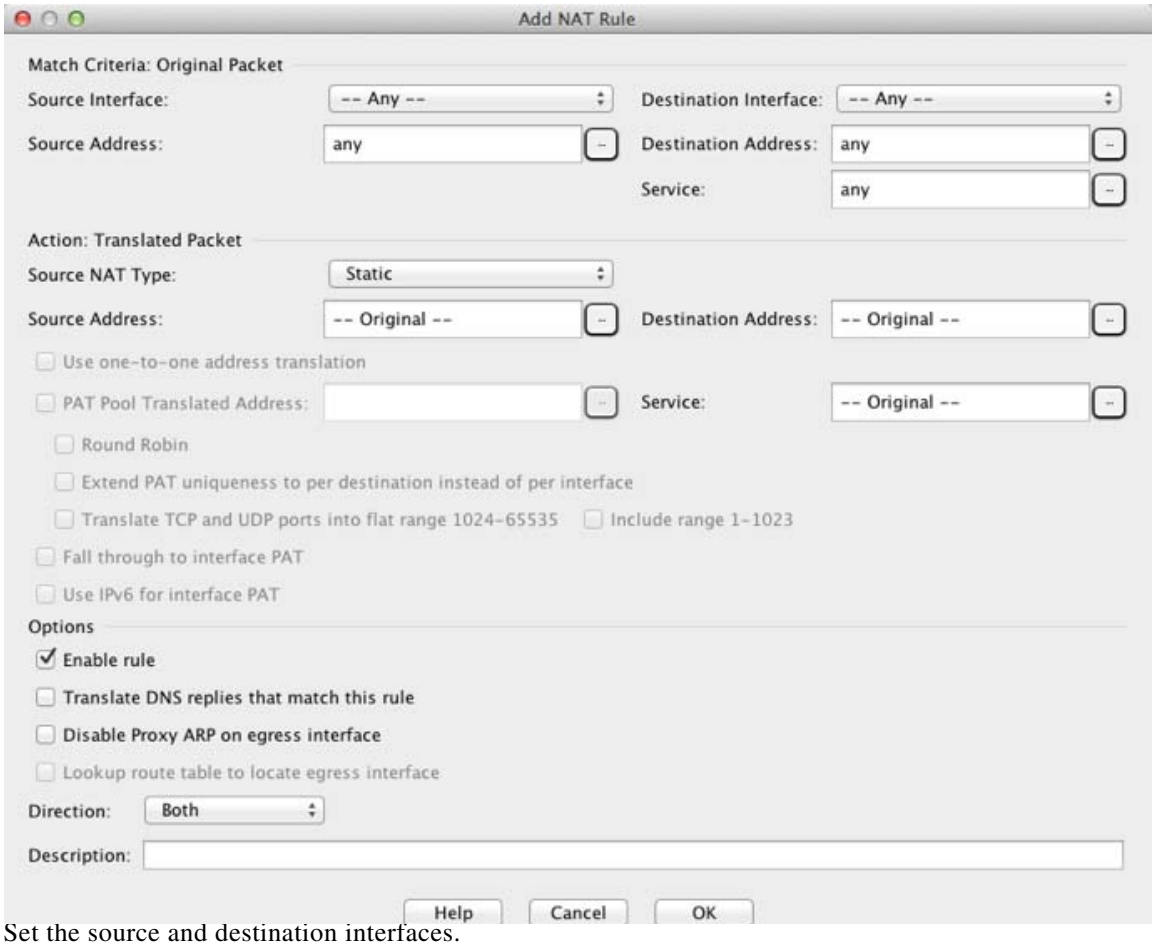

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

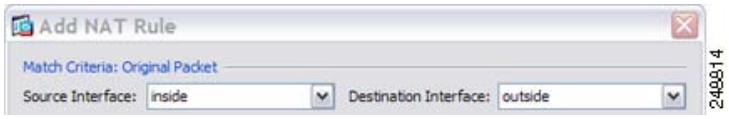

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

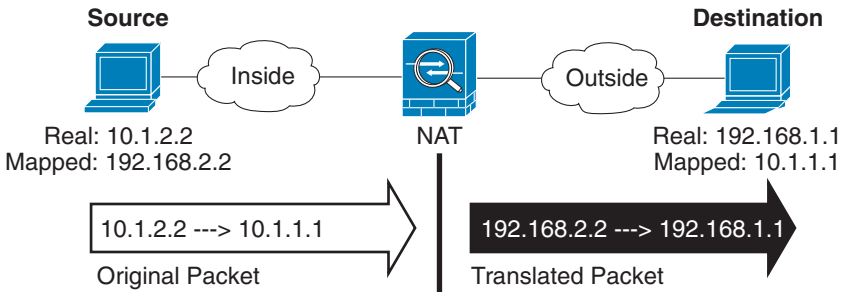

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

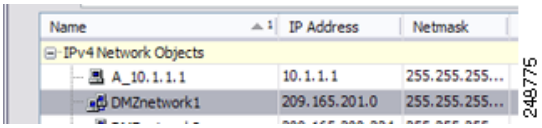

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see the "Main Differences Between Network Object NAT and Twice NAT" section on page 39-15.

Step 4 (Optional) Identify the original packet port (the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

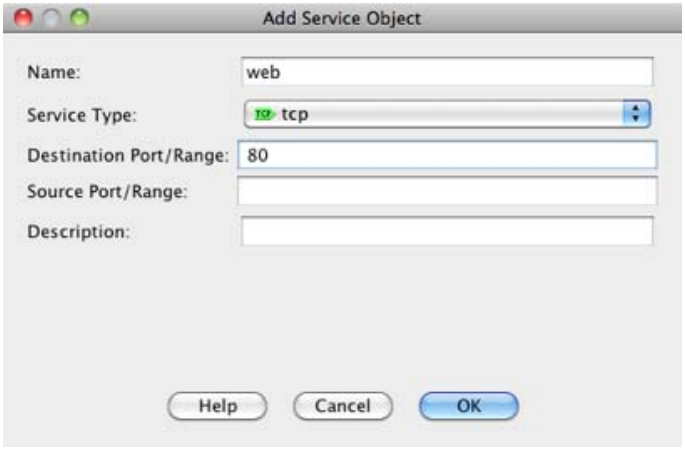

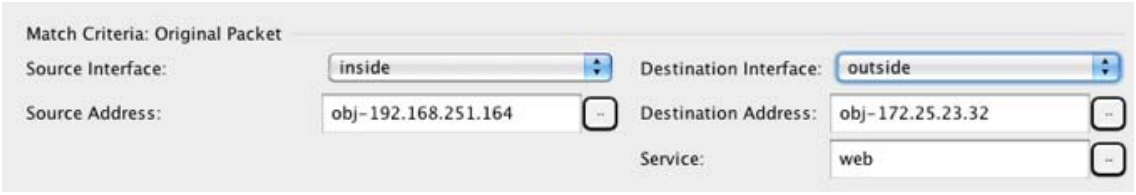

Step 5 Choose **Dynamic** from the Match Criteria: Translated Packet > Source NAT Type drop-down list.

This setting only applies to the source address; the destination translation is always static.

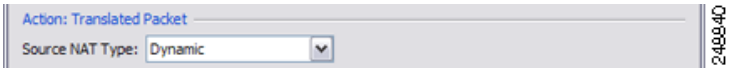

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original packet vs. the translated packet.

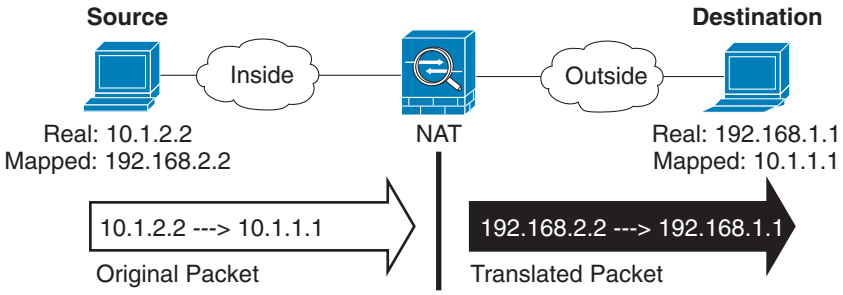

a. You can perform either dynamic NAT or Dynamic PAT using a PAT pool:

- Dynamic NAT—For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

  For dynamic NAT, you typically configure a larger group of source addresses to be mapped to a smaller group.

Note The object or group cannot contain a subnet.

- Dynamic PAT using a PAT pool—.To configure a PAT pool, check the **PAT Pool Translated Address** check box, then click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated PAT Pool Address dialog box. **Note**: Leave the Source Address field empty.

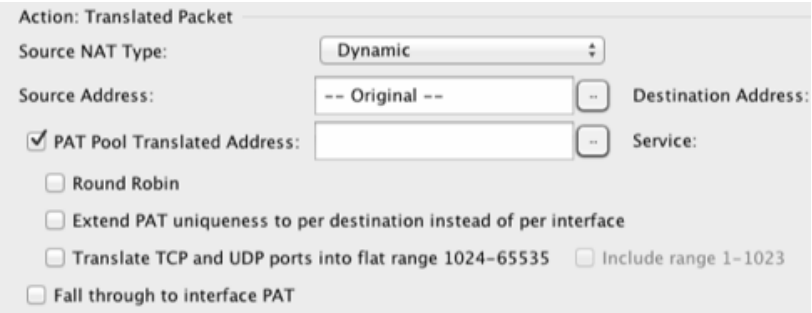

Note The object or group cannot contain a subnet.

(Optional) For a PAT pool, configure the following options:

  - To assign addresses/ports in a round-robin fashion, check the **Round Robin** check box. Without round-robin, by default, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns an address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
  - (8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Extend PAT uniqueness to per destination instead of per interface** check box to use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.
  - (8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Translate TCP or UDP ports into flat range (1024-65535)** check box to use the 1024 to 65535 port range as a single flat range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include range 1 to 1023** check box.

b. (Optional, Routed Mode Only) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.

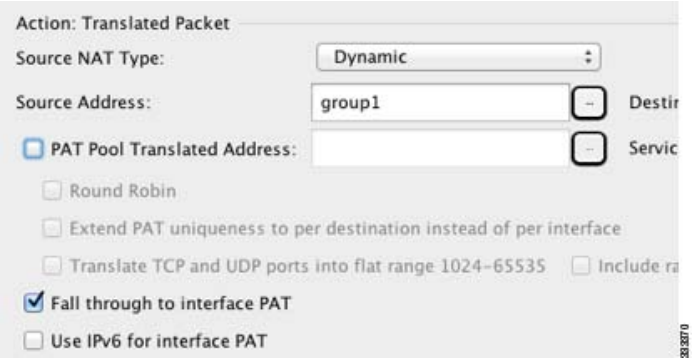

c. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see the "Static NAT" section on page 39-3. See the "Guidelines and Limitations" section on page 41-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation (see Step 7). For this option, you must configure a specific interface for the Source Interface in Step 2. See the "Static Interface NAT with Port Translation" section on page 39-6 for more information.

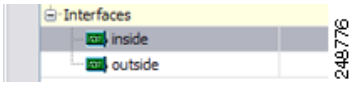

Step 7 (Optional) Identify the translated packet port (the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

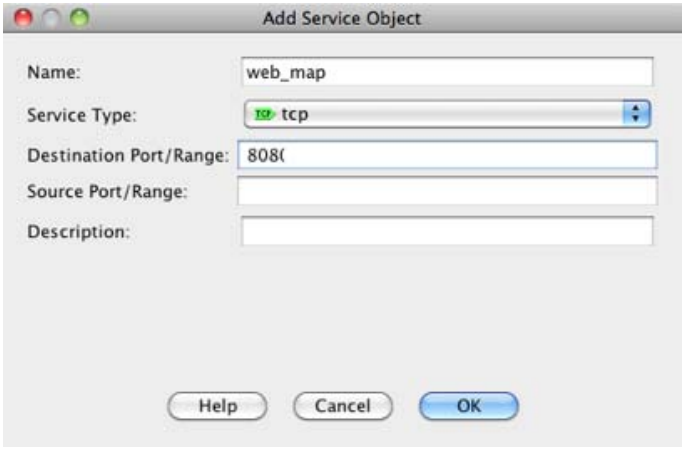

Step 8 (Optional) Configure NAT options in the Options area.

a. Enable rule —Enables this NAT rule. The rule is enabled by default.

b. (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See the "DNS and NAT" section on page 39-30 for more information.

c. Description—Adds a description about the rule up to 200 characters in length.

Step 9 Click **OK**.

Configuring Dynamic PAT (Hide)

This section describes how to configure twice NAT for dynamic PAT (hide). For dynamic PAT using a PAT pool, see the "Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool" section on page 41-4 instead of using this section. For more information, see the "Dynamic PAT" section on page 39-10.

Detailed Steps

To configure dynamic PAT, perform the following steps:

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

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

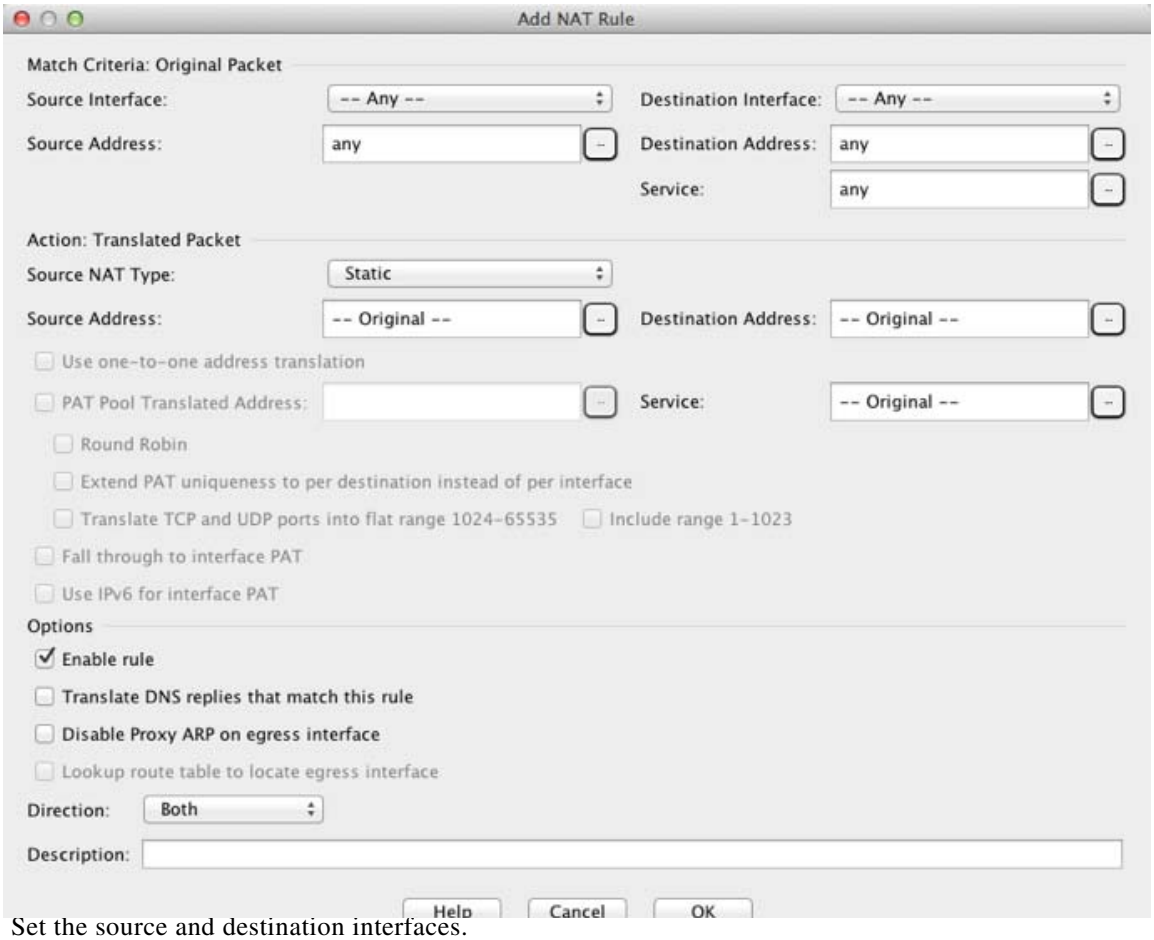

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

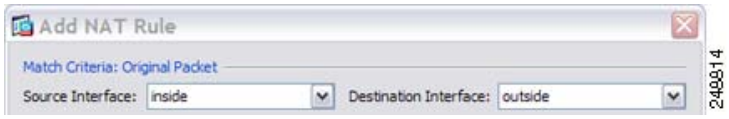

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

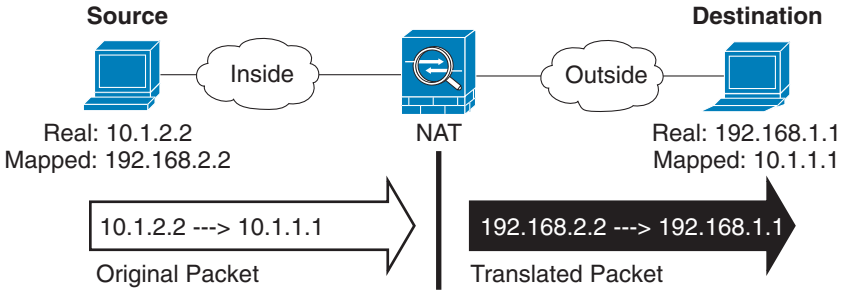

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

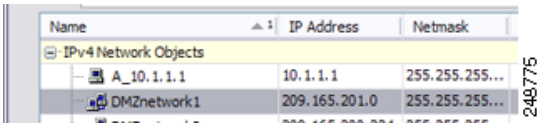

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see the "Main Differences Between Network Object NAT and Twice NAT" section on page 39-15.

Step 4 (Optional) Identify the original packet port (the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

Dynamic PAT does not support additional port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

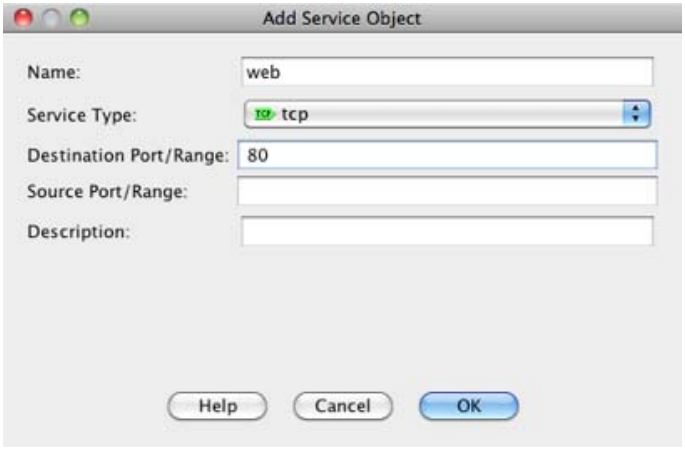

Step 5 Choose **Dynamic PAT (Hide)** from the Match Criteria: Translated Packet > Source NAT Type drop-down list.

This setting only applies to the source address; the destination translation is always static.

Note To configure dynamic PAT using a PAT pool, choose **Dynamic** instead of Dynamic PAT (Hide), see the "Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool" section on page 41-4.

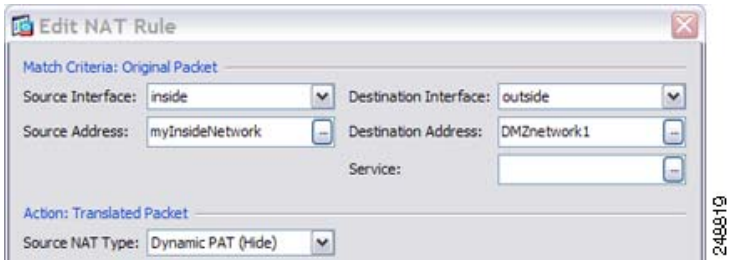

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original packet vs. the translated packet.

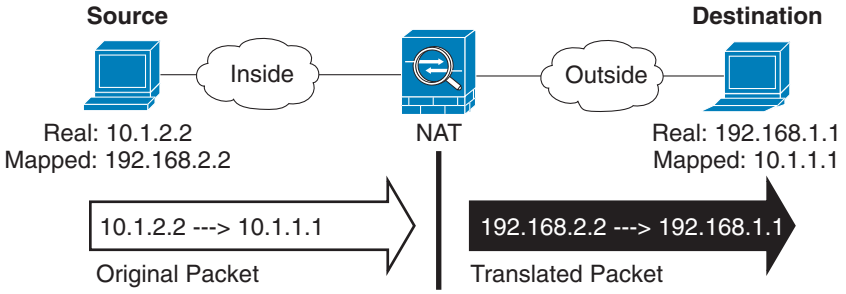

a. For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an existing network object or interface or create a new object from the Browse Translated Source Address dialog box.

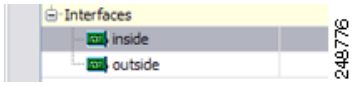

If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.

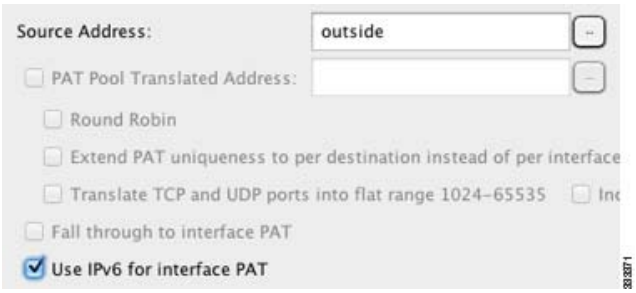

b. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see the "Static NAT" section on page 39-3. See the "Guidelines and Limitations" section on page 41-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation (see Step 7). For this option, you must configure a specific interface for the Source Interface in Step 2. See the "Static Interface NAT with Port Translation" section on page 39-6 for more information.

Step 7 (Optional) Identify the translated packet port (the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object from the Browse Translated Service dialog box.

You can also create a new service object from the Browse Translated Service dialog box and use this object as the mapped destination port.

Dynamic PAT does not support additional port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

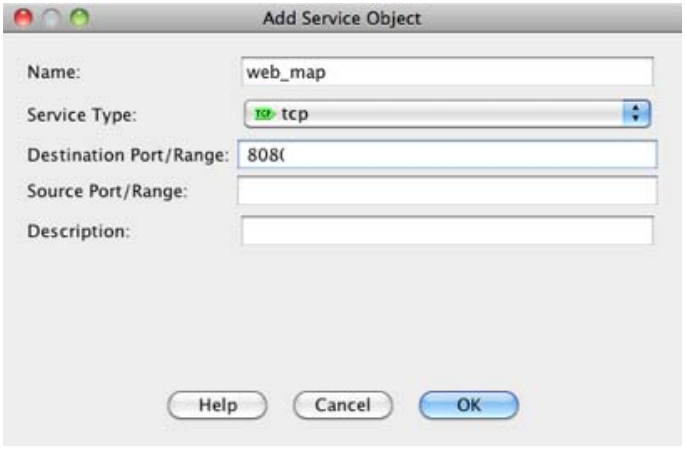

Step 8 (Optional) Configure NAT options in the Options area.

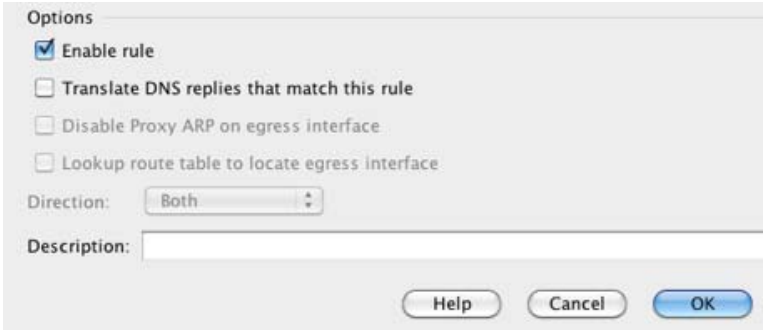

a. Enable rule —Enables this NAT rule. The rule is enabled by default.

b. (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See the "DNS and NAT" section on page 39-30 for more information.

c. Description—Adds a description about the rule up to 200 characters in length.

Step 9 Click **OK**.

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

This section describes how to configure a static NAT rule using twice NAT. For more information about static NAT, see the "Static NAT" section on page 39-3.

Detailed Steps

To configure static NAT, perform the following steps:

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

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

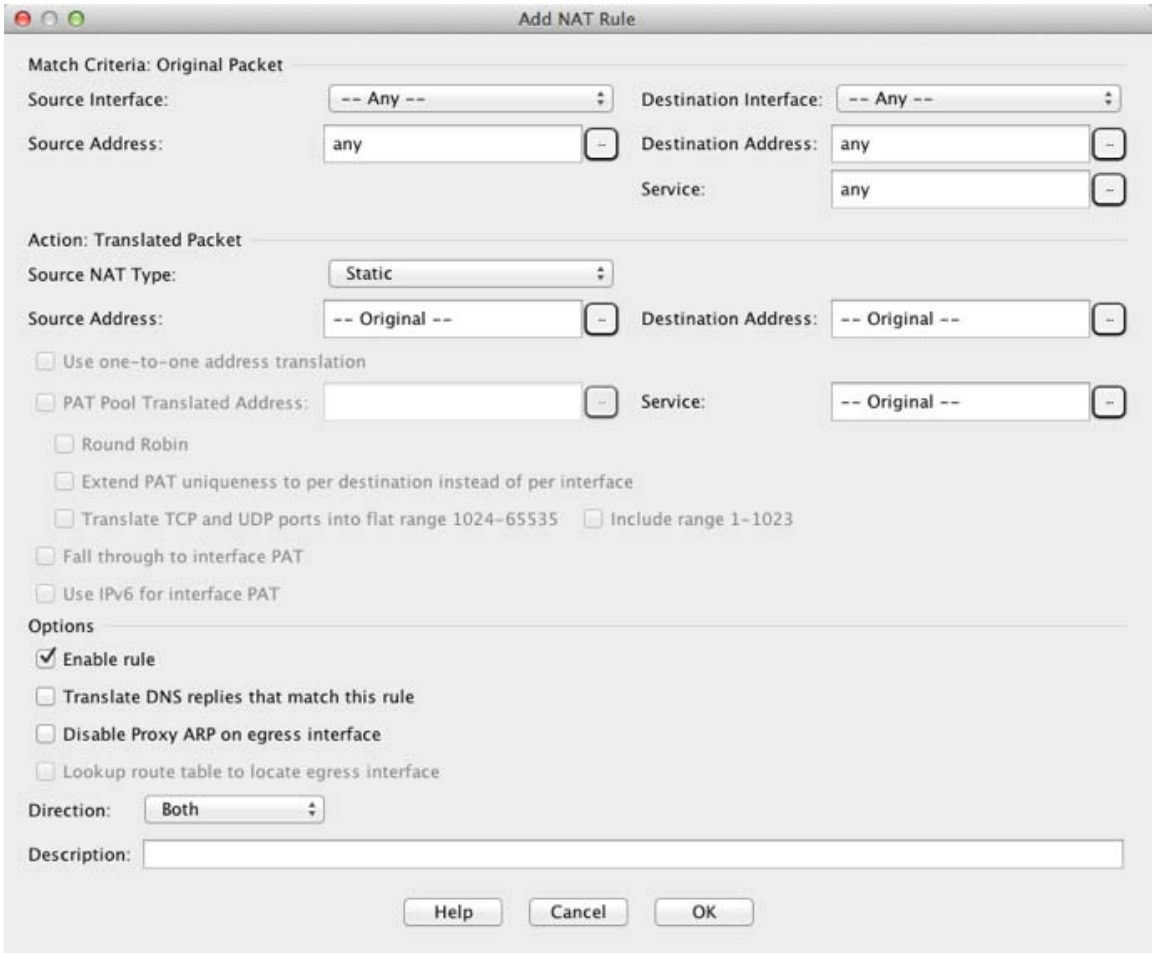

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

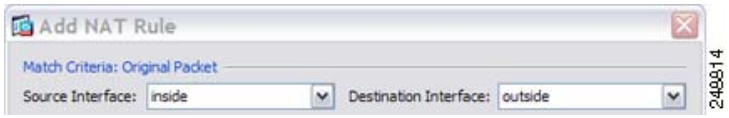

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

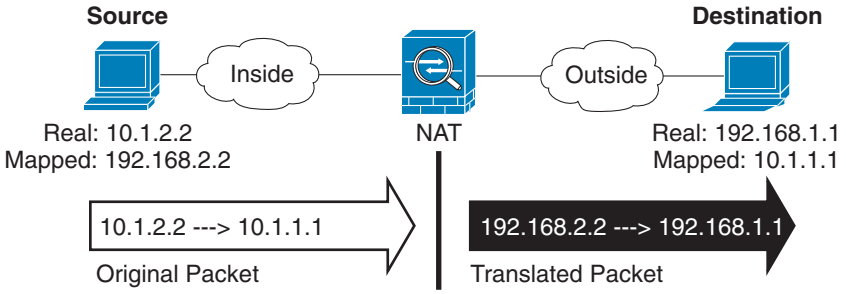

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**, but do not use this option except for identity NAT. See the "Configuring Identity NAT" section on page 41-24 for more information.

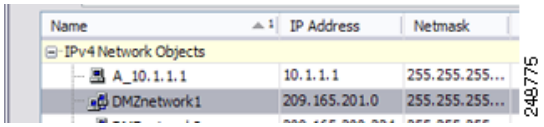

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see the "Main Differences Between Network Object NAT and Twice NAT" section on page 39-15.

Step 4 (Optional) Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

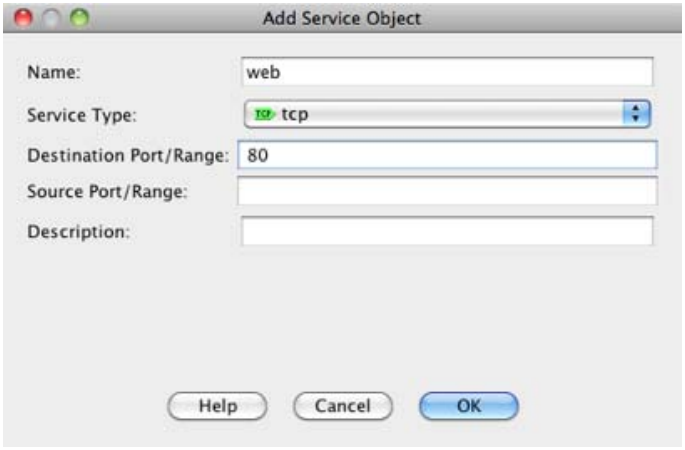

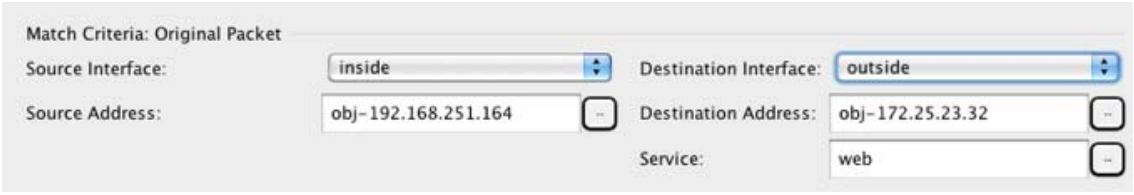

Step 5 Choose **Static** from the Match Criteria: Translated Packet > Source NAT Type drop-down list. Static is the default setting.

This setting only applies to the source address; the destination translation is always static.

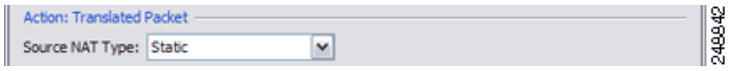

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original packet vs. the translated packet.

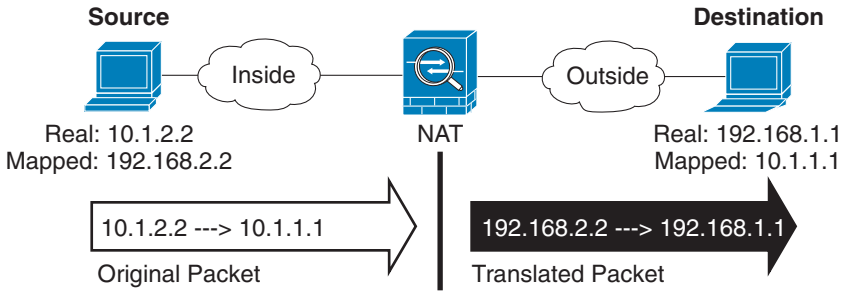

a. For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

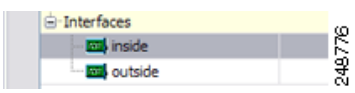

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.

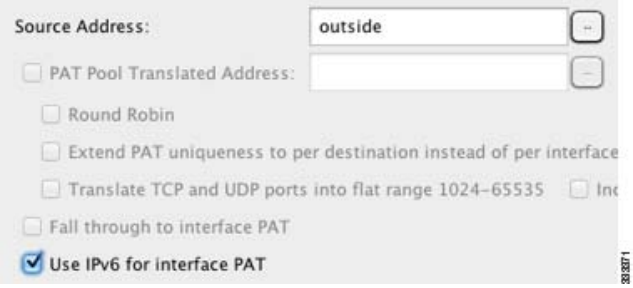

For more information, see the "Static Interface NAT with Port Translation" section on page 39-6. See the "Guidelines and Limitations" section on page 41-2 for information about disallowed mapped IP addresses.

b. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. For more information, see the "Static Interface NAT with Port Translation" section on page 39-6. See the "Guidelines and Limitations" section on page 41-2 for information about disallowed mapped IP addresses.

Step 7 (Optional) Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

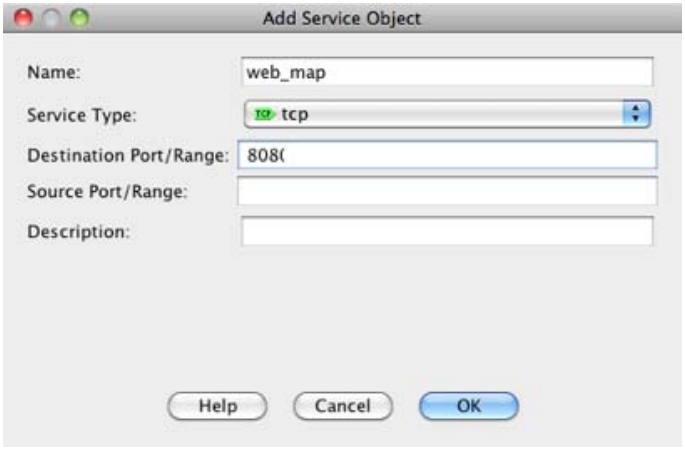

Step 8 (Optional) For NAT46, check the **Use one-to-one address translation** check box. For NAT46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 9 (Optional) Configure NAT options in the Options area.

a. Enable rule —Enables this NAT rule. The rule is enabled by default.

b. (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See the "DNS and NAT" section on page 39-30 for more information.

c. Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See the "Mapped Addresses and Routing" section on page 39-22 for more information.

d. Direction—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses.

e. Description—Adds a description about the rule up to 200 characters in length.

Step 10 Click **OK**.

Configuring Identity NAT

This section describes how to configure an identity NAT rule using twice NAT. For more information about identity NAT, see the "Identity NAT" section on page 39-12.

Detailed Steps

To configure identity NAT, perform the following steps:

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

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

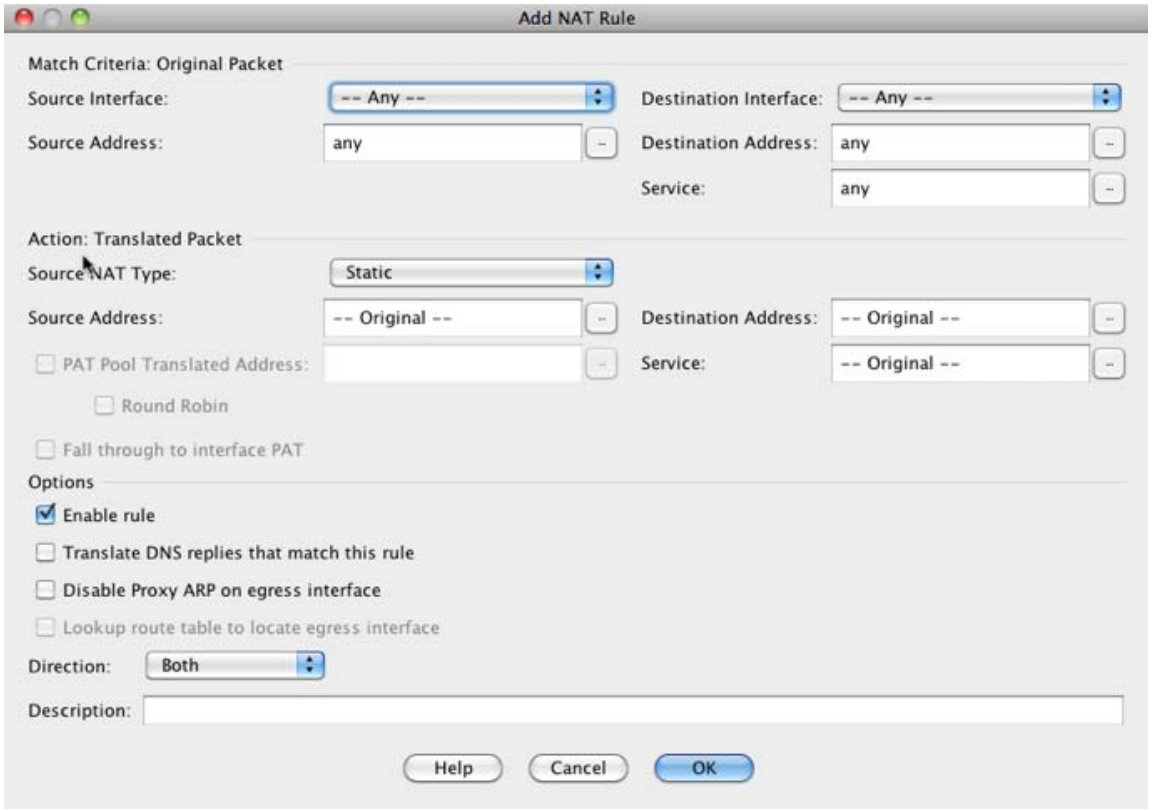

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

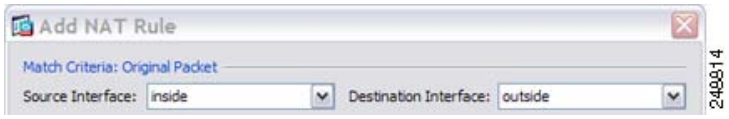

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.

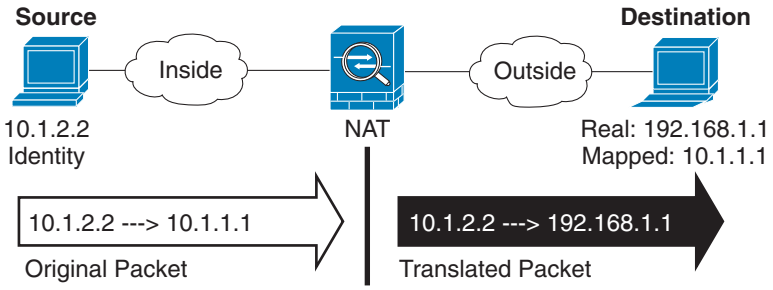

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**; only use this option when also setting the mapped address to **any**.

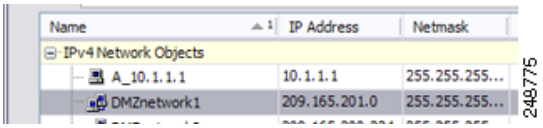

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see the "Main Differences Between Network Object NAT and Twice NAT" section on page 39-15.

Step 4 (Optional) Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

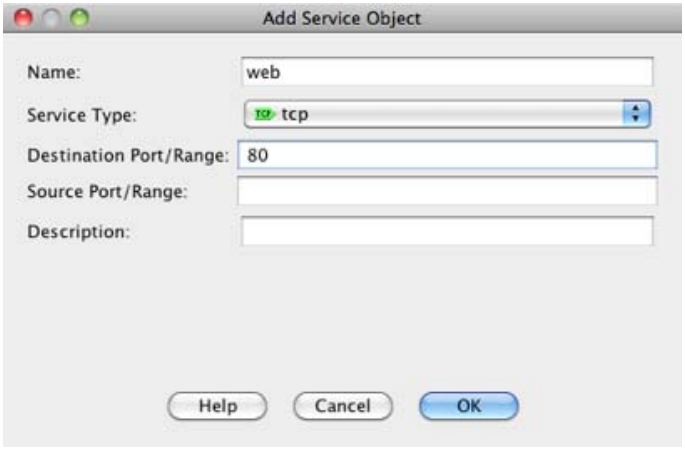

Step 5 Choose **Static** from the Match Criteria: Translated Packet > Source NAT Type drop-down list. Static is the default setting.

This setting only applies to the source address; the destination translation is always static.

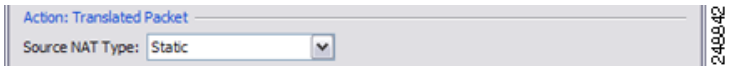

Step 6 Identify the translated packet addresses; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.

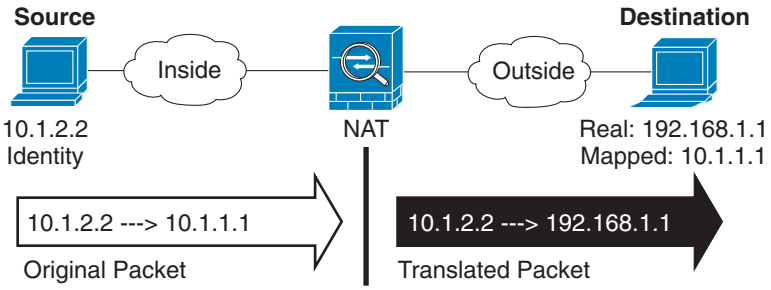

a. For the Match Criteria: Translated Packet > Source Address, click the browse button and choose the same network object or group from the Browse Translated Source Address dialog box that you chose for the real source address. Use **any** if you specified **any** for the real address.

b. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see the "Static NAT" section on page 39-3. See the "Guidelines and Limitations" section on page 41-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface. If you specify an interface, be sure to also configure a a service translation. For more information, see the "Static Interface NAT with Port Translation" section on page 39-6.

Step 7 (Optional) Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real destination port. NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports. The "not equal" (**!=**) operator is not supported.

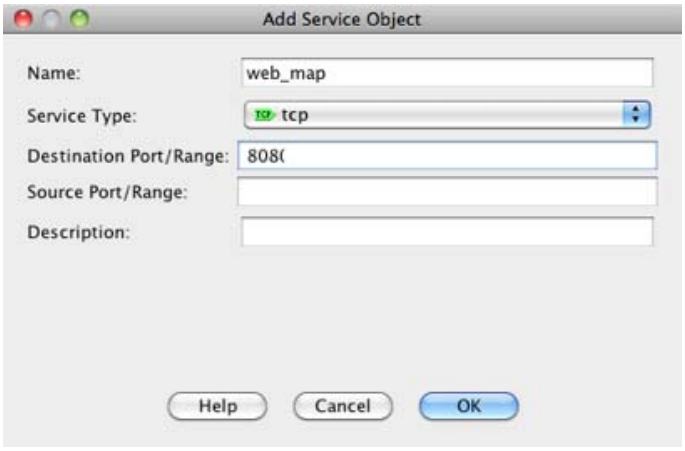

Step 8 (Optional) Configure NAT options in the Options area.

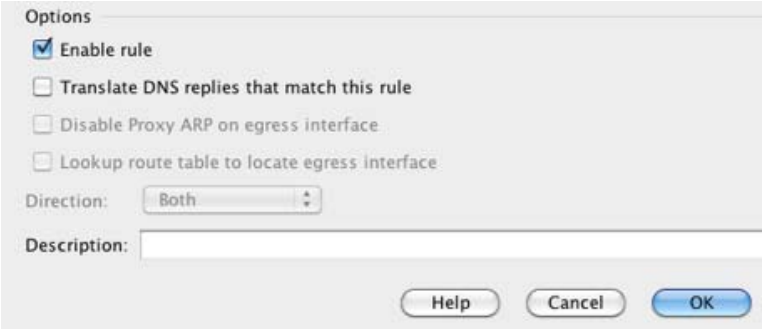

a. Enable rule —Enables this NAT rule. The rule is enabled by default.

b. Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See the "Mapped Addresses and Routing" section on page 39-22 for more information.

c. (Routed mode; interface(s) specified) Lookup route table to locate egress interface—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See the "Determining the Egress Interface" section on page 39-24 for more information.

d. Direction—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses. You might want to use this setting for testing purposes.

e. Description—Adds a description about the rule up to 200 characters in length.

Note Although the "Translate DNS replies that match this rule" check box is available if you do not configure a destination address, this option is not applicable to identity NAT because you are translating the address to itself, so the DNS reply does not need modification. See the "DNS and NAT" section on page 39-30 for more information.

Step 9 Click **OK**.

Configuring Per-Session PAT Rules

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT. For more information about per-session vs. multi-session PAT, see the "Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later)" section on page 39-11.

Detailed Steps

To configure a per-session PAT rule, see the "Configuring Per-Session PAT Rules" section on page 40-18.

Monitoring Twice NAT

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

Fields

- Available Graphs—Lists the components you can graph.
  - Xlate Utilization—Displays the ASA NAT utilization.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected entry from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

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

Fields

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

Configuration Examples for Twice NAT

This section includes the following configuration examples:

- Different Translation Depending on the Destination (Dynamic PAT), page 41-31
- Different Translation Depending on the Destination Address and Port (Dynamic PAT), page 41-39

Different Translation Depending on the Destination (Dynamic PAT)

Figure 41-1 shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 41-1 Twice NAT with Different Destination Addresses

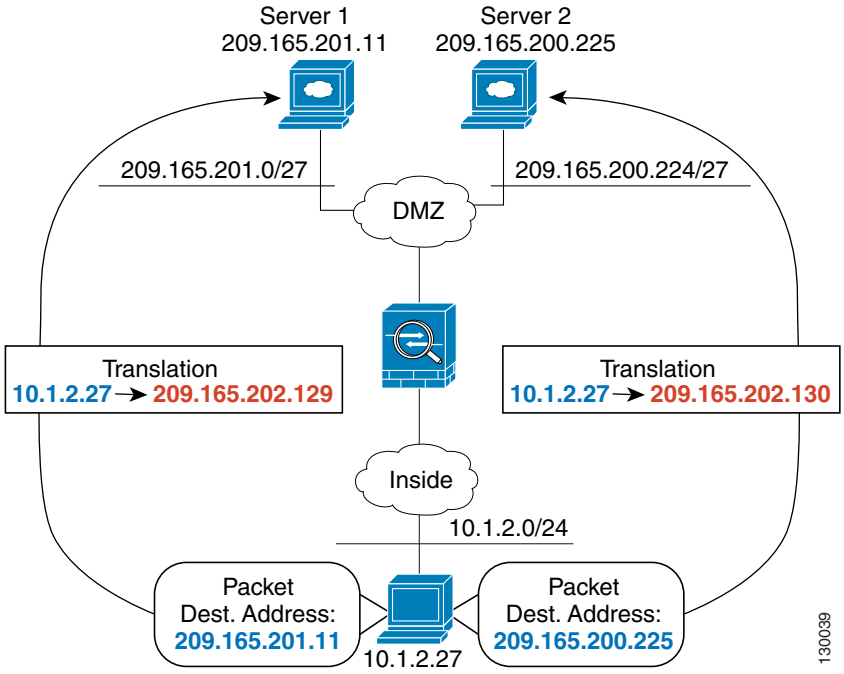

Step 1 Add a NAT rule for traffic from the inside network to DMZ network 1:

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

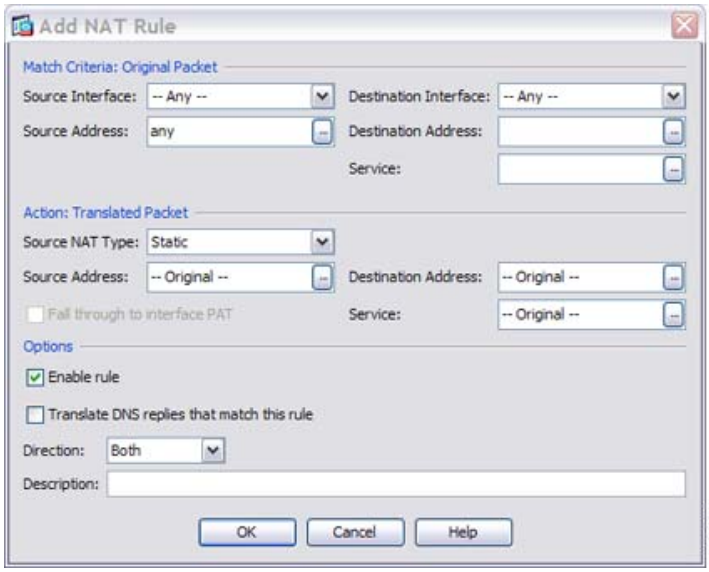

Step 2 Set the source and destination interfaces:

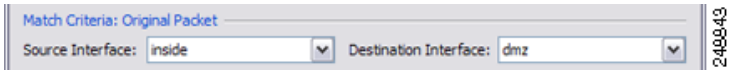

Step 3 For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

a. Add the new network object.

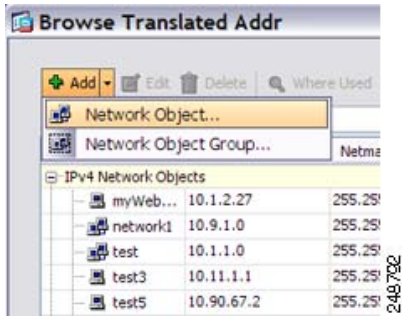

b. Define the inside network addresses, and click **OK**.

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

Step 4 For the Original Destination Address, click the browse button to add a new network object for DMZ network 1 in the Browse Original Destination Address dialog box.

a. Add the new network object.

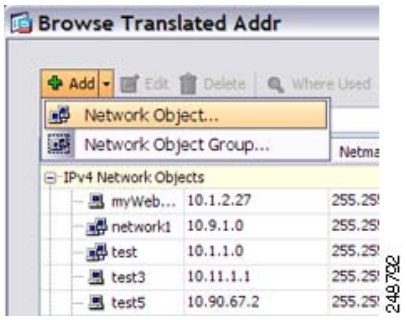

b. Define the DMZ network 1 addresses, and click **OK**.

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

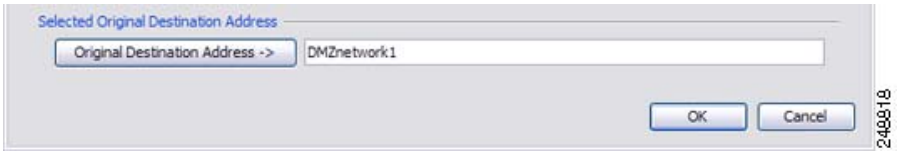

Step 5 Set the NAT Type to **Dynamic PAT (Hide)**:

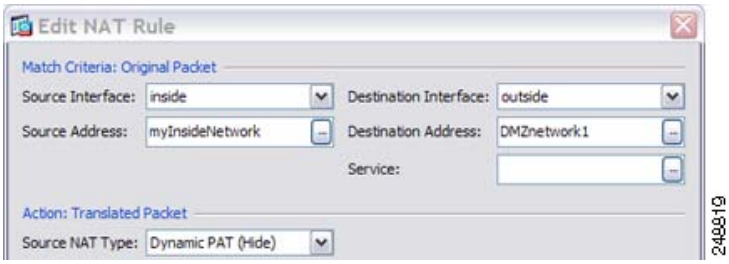

Step 6 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.

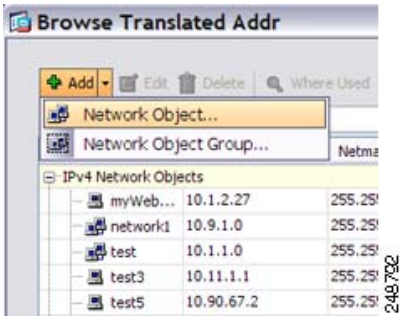

b. Define the PAT address, and click **OK**.

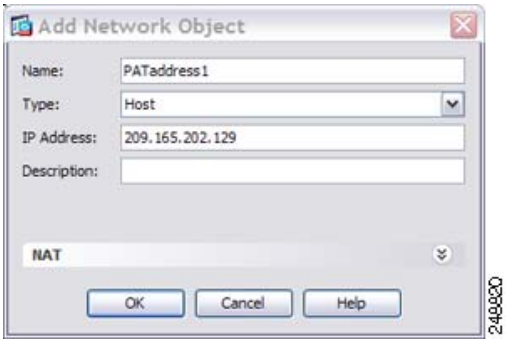

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

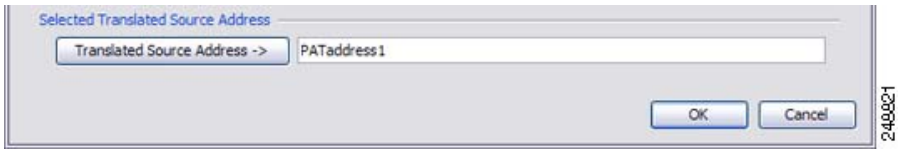

Step 7 For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork1) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.

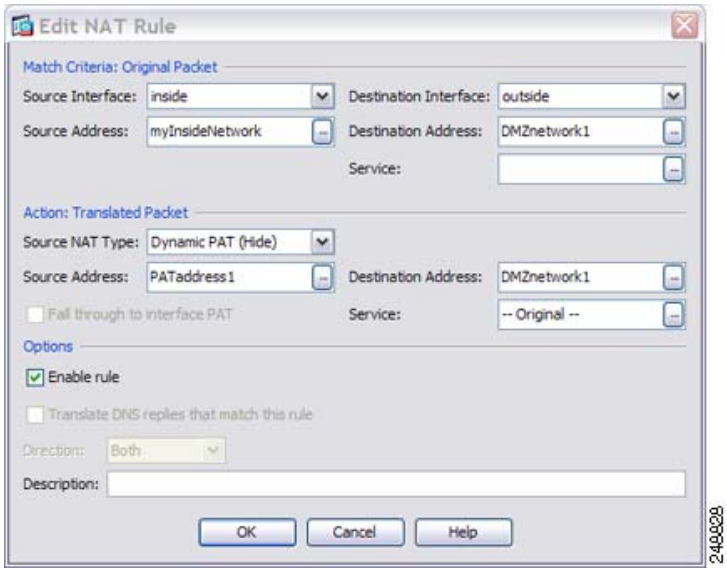

Step 8 Click **OK** to add the rule to the NAT table.

Step 9 Add a NAT rule for traffic from the inside network to DMZ network 2:

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

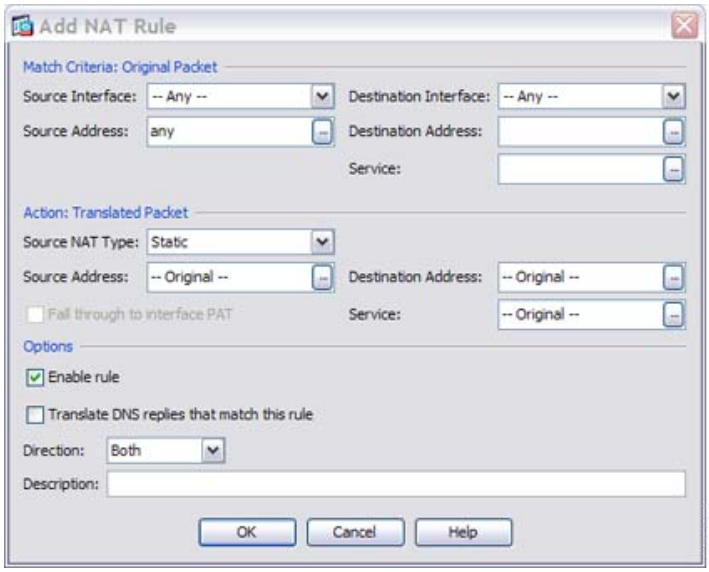

Step 10 Set the source and destination interfaces:

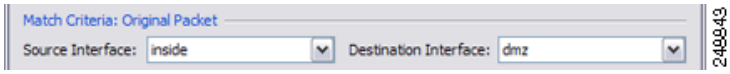

Step 11 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 12 For the Original Destination Address, click the browse button to add a new network object for DMZ network 2 in the Browse Original Destination Address dialog box.

a. Add the new network object.

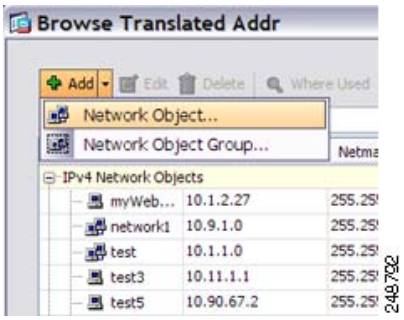

b. Define the DMZ network 2 addresses, and click **OK**.

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

Step 13 Set the NAT Type to **Dynamic PAT (Hide)**:

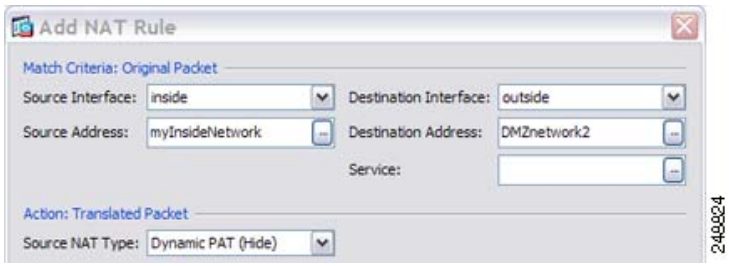

Step 14 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.

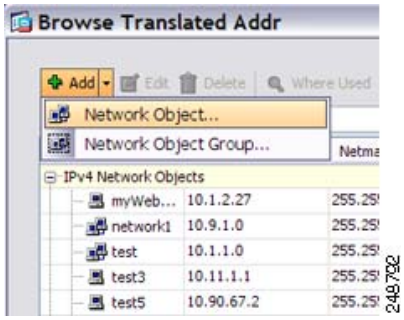

b. Define the PAT address, and click **OK**.

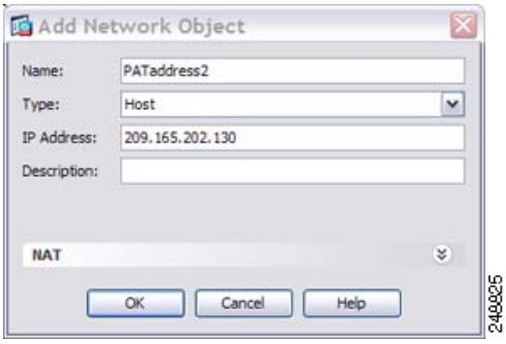

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

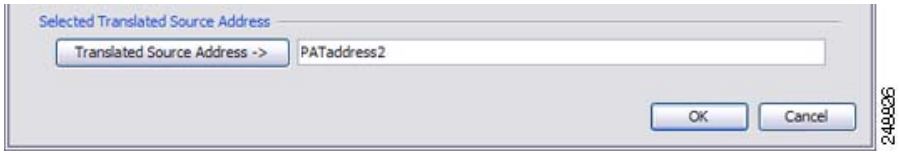

Step 15 For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork2) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.

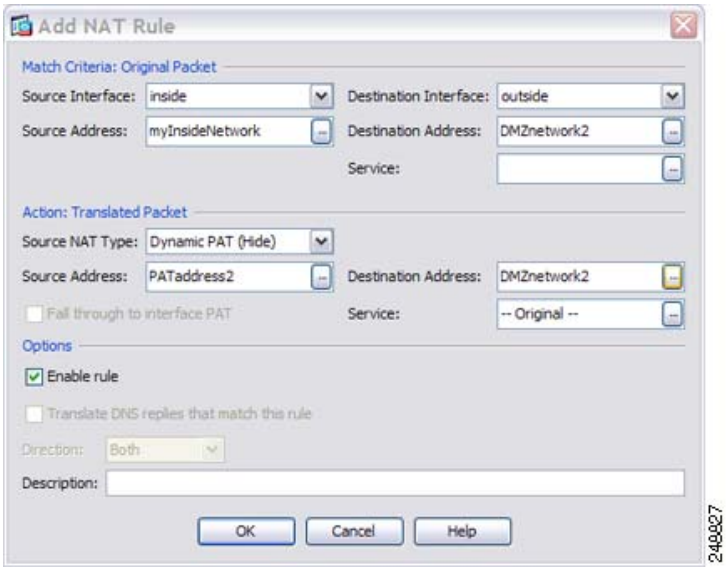

Step 16 Click **OK** to add the rule to the NAT table.

Step 17 Click **Apply**.

Different Translation Depending on the Destination Address and Port (Dynamic PAT)

Figure 41-2 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129:*port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130:*port*.

Figure 41-2 Twice NAT with Different Destination Ports

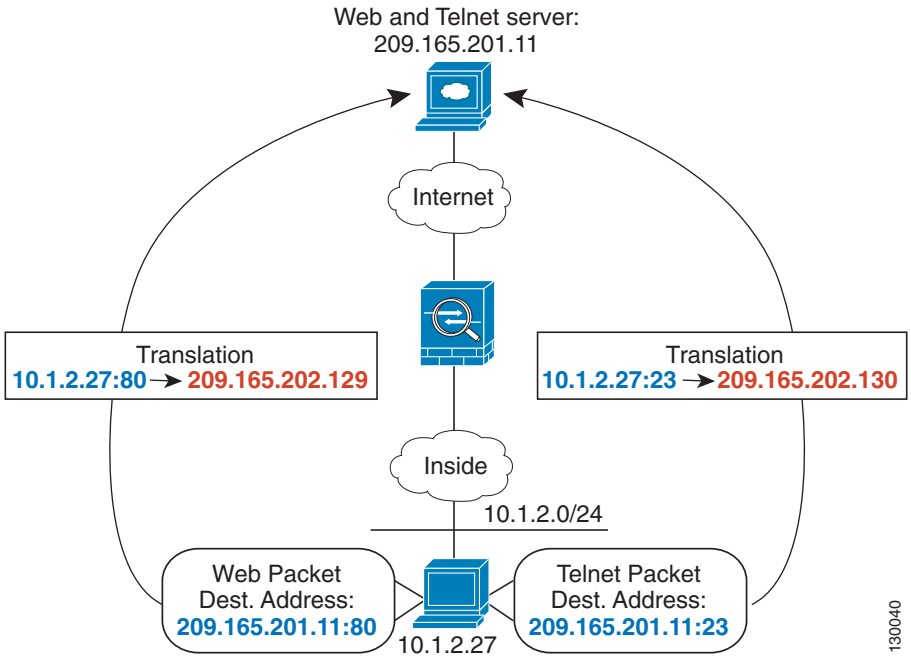

Step 1 Add a NAT rule for traffic from the inside network to the Telnet server:

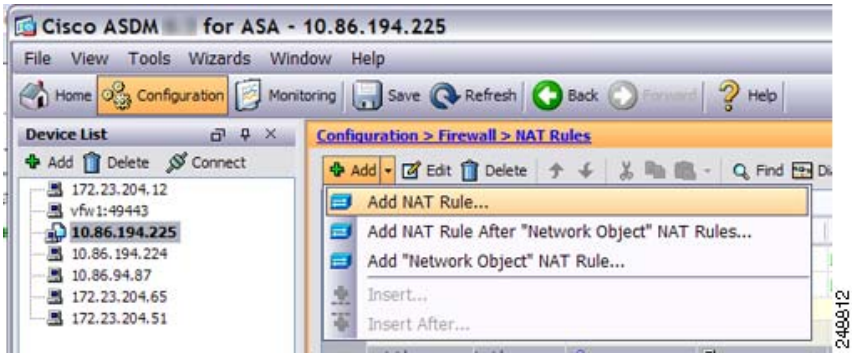

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

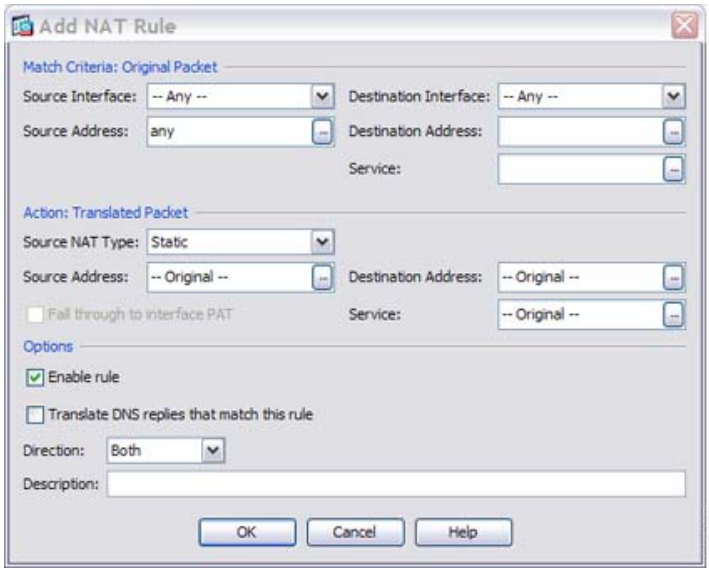

Step 2 Set the source and destination interfaces:

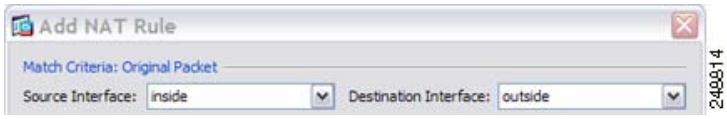

Step 3 For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

a. Add the new network object.

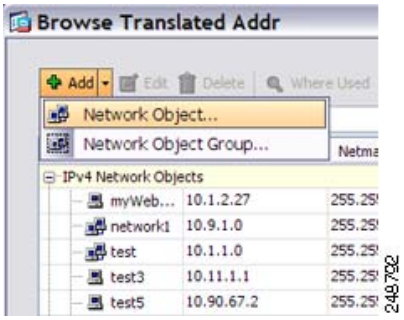

b. Define the inside network addresses, and click **OK**.

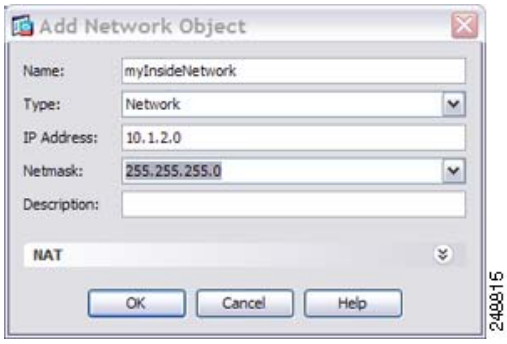

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 4 For the Original Destination Address, click the browse button to add a new network object for the Telnet/Web server in the Browse Original Destination Address dialog box.

a. Add the new network object.

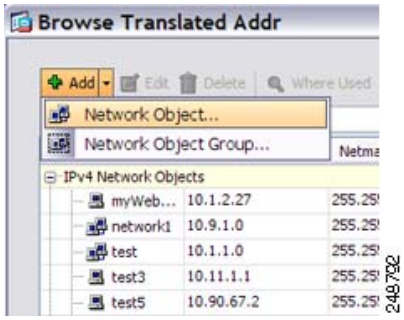

b. Define the server address, and click **OK**.

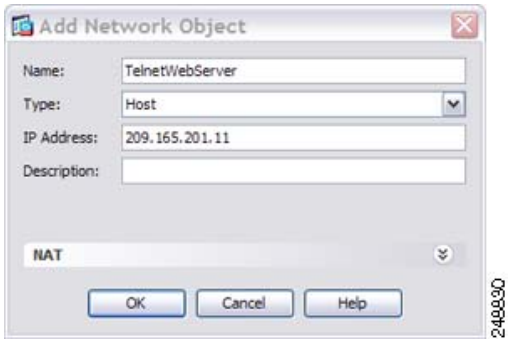

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5 For the Original Service, click the browse button to add a new service object for Telnet in the Browse Original Service dialog box.

a. Add the new service object.

b. Define the protocol and port, and click **OK**.

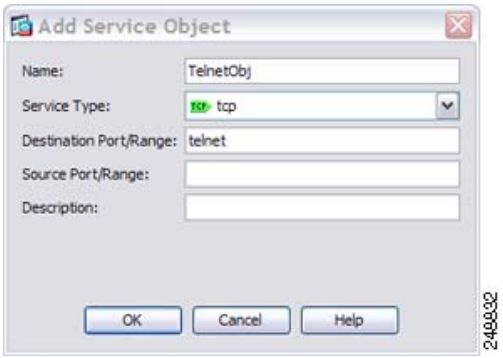

c. Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 6 Set the NAT Type to **Dynamic PAT (Hide)**:

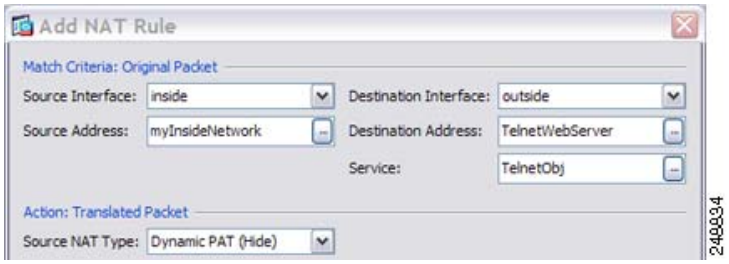

Step 7 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.

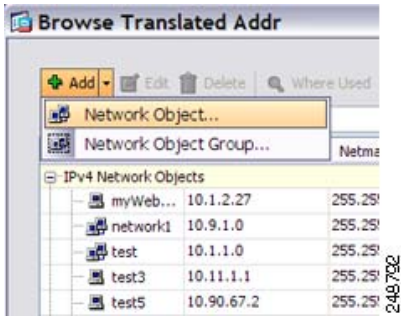

b. Define the PAT address, and click **OK**.

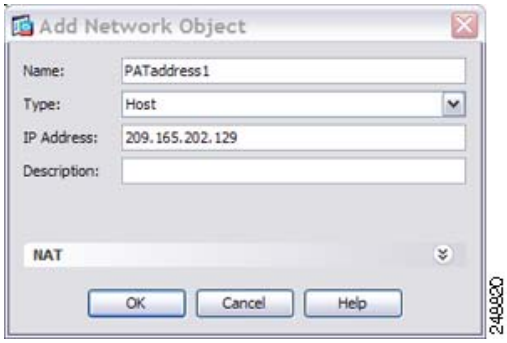

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

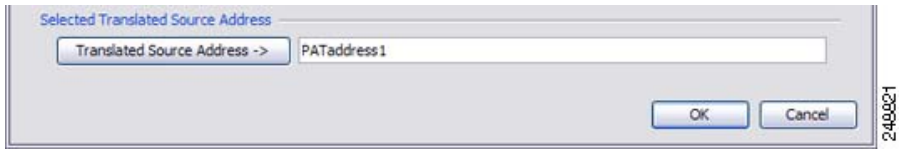

Step 8 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.

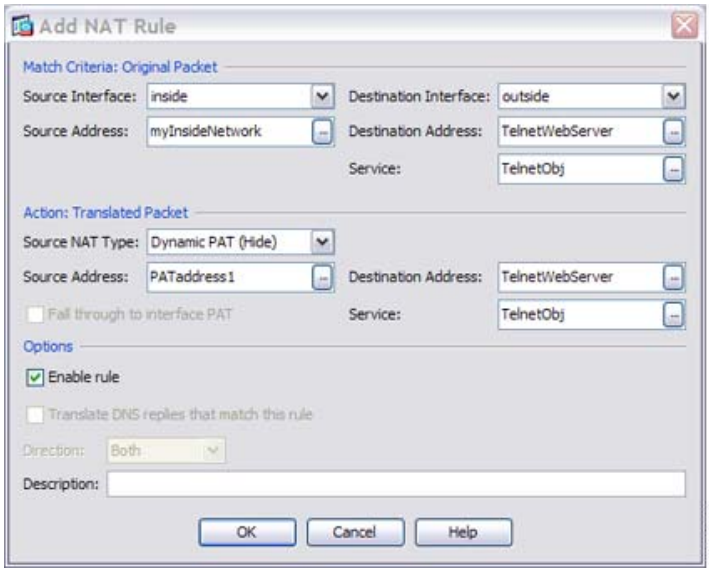

Step 9 Click **OK** to add the rule to the NAT table.

Step 10 Add a NAT rule for traffic from the inside network to the web server:

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

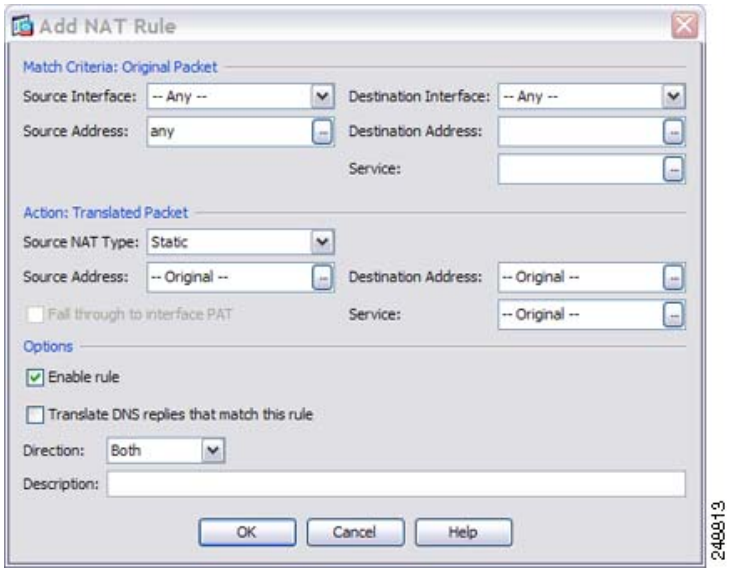

Step 11 Set the real and mapped interfaces:

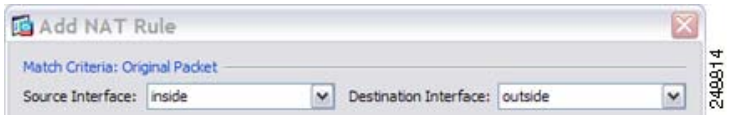

Step 12 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 13 For the Original Destination Address, type the name of the Telnet/web server network object (TelnetWebServer) or click the browse button to choose it.

Step 14 For the Original Service, click the browse button to add a new service object for HTTP in the Browse Original Service dialog box.

a. Add the new service object.

b. Define the protocol and port, and click **OK**.

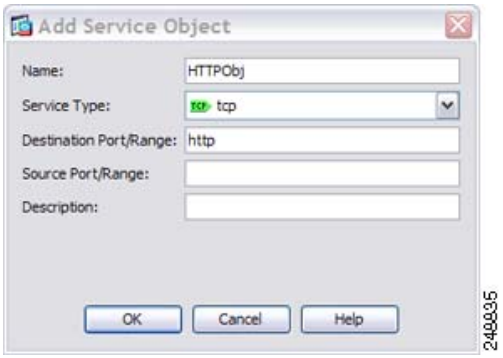

c. Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 15 Set the NAT Type to **Dynamic PAT (Hide)**:

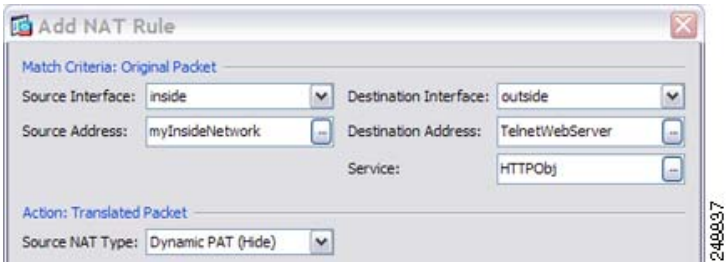

Step 16 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.

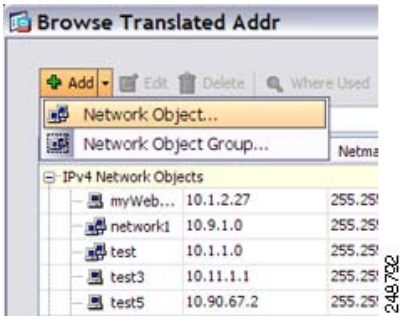

b. Define the PAT address, and click **OK**.

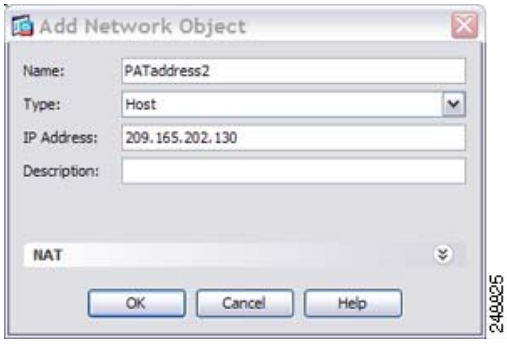

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

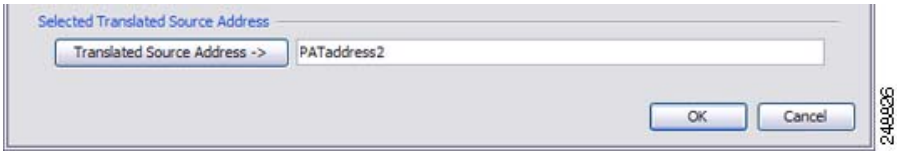

Step 17 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the Original and Translated destination addresses.

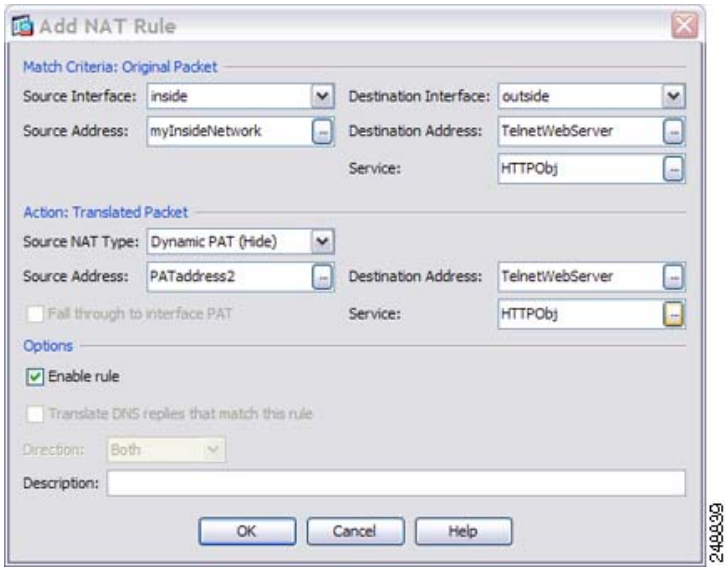

Step 18 Click **OK** to add the rule to the NAT table.

Step 19 Click **Apply**.

Feature History for Twice NAT

Table 41-1 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 41-1 Feature History for Twice NAT

Feature Name	Platform Releases	Feature Information
Twice NAT	8.3(1)	Twice NAT lets you identify both the source and destination address in a single rule. We modified the following screen: Configuration > Firewall > NAT Rules.
Identity NAT configurable proxy ARP and route lookup	8.4(2)/8.5(1)	In earlier releases for identity NAT, proxy ARP was disabled, and a route lookup was always used to determine the egress interface. You could not configure these settings. In 8.4(2) and later, the default behavior for identity NAT was changed to match the behavior of other static NAT configurations: proxy ARP is enabled, and the NAT configuration determines the egress interface (if specified) by default. You can leave these settings as is, or you can enable or disable them discretely. Note that you can now also disable proxy ARP for regular static NAT. For pre-8.3 configurations, the migration of NAT exempt rules (the **nat 0 access-list** command) to 8.4(2) and later now includes the following keywords to disable proxy ARP and to use a route lookup: **no-proxy-arp** and **route-lookup**. The **unidirectional** keyword that was used for migrating to 8.3(2) and 8.4(1) is no longer used for migration. When upgrading to 8.4(2) from 8.3(1), 8.3(2), and 8.4(1), all identity NAT configurations will now include the **no-proxy-arp** and **route-lookup** keywords, to maintain existing functionality. The **unidirectional** keyword is removed. We modified the following screen: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule
PAT pool and round robin address assignment	8.4(2)/8.5(1)	You can now specify a pool of PAT addresses instead of a single address. You can also optionally enable round-robin assignment of PAT addresses instead of first using all ports on a PAT address before using the next address in the pool. These features help prevent a large number of connections from a single PAT address from appearing to be part of a DoS attack and makes configuration of large numbers of PAT addresses easy. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.

Table 41-1 Feature History for Twice NAT (continued)

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

Table 41-1 Feature History for Twice NAT (continued)

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

Table 41-1 Feature History for Twice NAT (continued)

Feature Name	Platform Releases	Feature Information
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode. We modified the following screen: Configuration > Firewall > NAT Rules.
Per-session PAT	9.0(1)	The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.

Configuring NAT (ASA 8.2 and Earlier)

This chapter describes Network Address Translation, and includes the following sections:

- NAT Overview, page 42-1
- Configuring NAT Control, page 42-15
- Using Dynamic NAT, page 42-16
- Using Static NAT, page 42-26
- Using NAT Exemption, page 42-32

NAT Overview

This section describes how NAT works on the ASA, and includes the following topics:

- Introduction to NAT, page 42-1
- NAT in Routed Mode, page 42-2
- NAT in Transparent Mode, page 42-3
- NAT Control, page 42-4
- NAT Types, page 42-6
- Policy NAT, page 42-10
- NAT and Same Security Level Interfaces, page 42-12
- Order of NAT Rules Used to Match Real Addresses, page 42-13
- Mapped Address Guidelines, page 42-13
- DNS and NAT, page 42-13

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 ASA 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 "Security Levels" section on page 17-1 for more information about security levels. See the "NAT Control" section on page 42-4 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 57-1 on page 57-4 for information about protocols that do not support NAT.

NAT in Routed Mode

Figure 42-1 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 42-1 NAT Example: Routed Mode

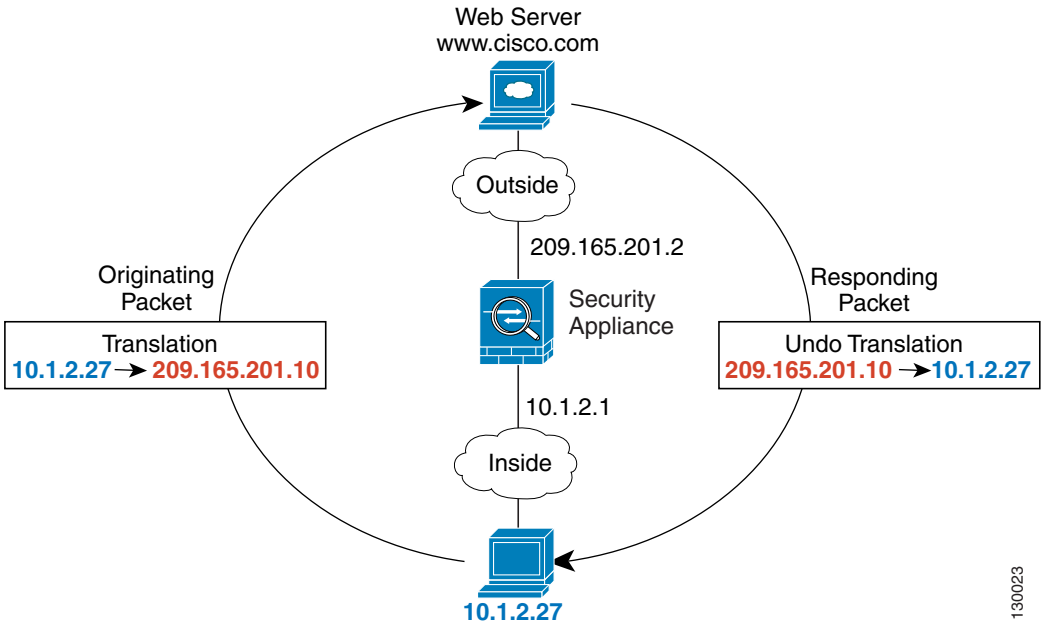

NAT in Transparent Mode

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 ASA 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 ASA).
- When you have VoIP or DNS traffic with NAT and inspection enabled, to successfully translate the IP address inside VoIP and DNS packets, the ASA needs to perform a route lookup. Unless the host is on a directly-connected network, then you need to add a static route on the ASA for the real host address that is embedded in the packet.
- 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 42-2 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 ASA receives the packet because the upstream router includes this mapped network in a static route directed through the ASA. The ASA 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 ASA sends it directly to the host.

Figure 42-2 NAT Example: Transparent Mode

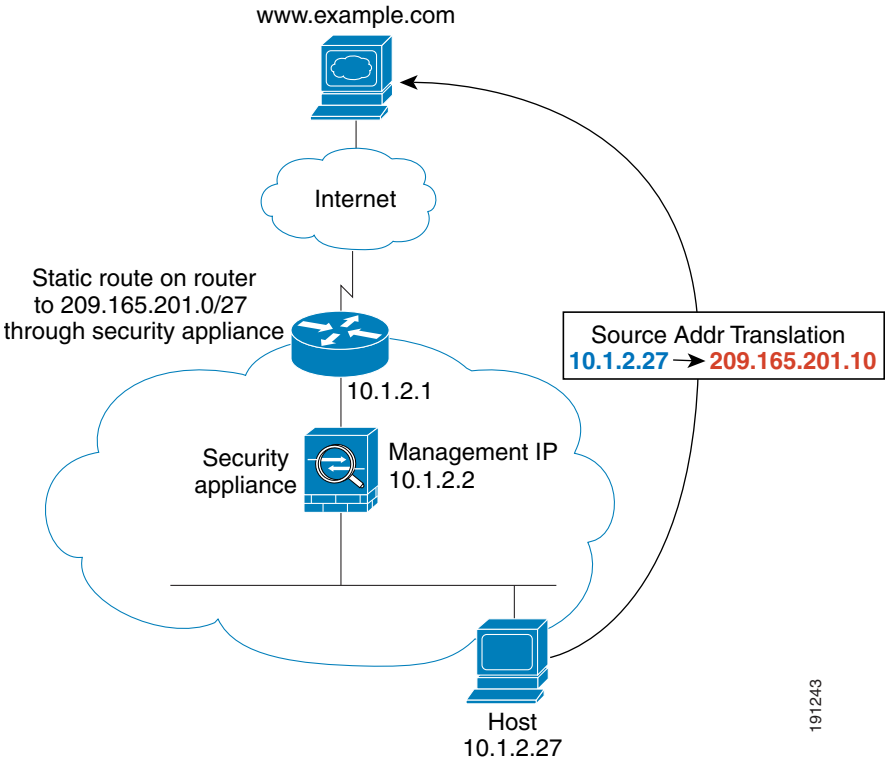

NAT Control

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 42-3.

Figure 42-3 NAT Control and Outbound Traffic

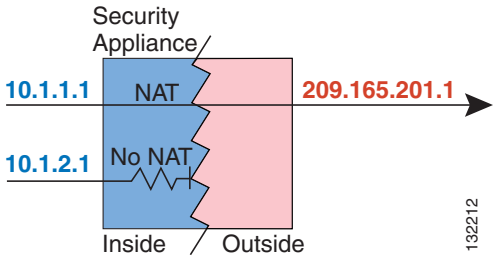

Interfaces at the same security level are not required to use NAT to communicate. However, if you configure dynamic NAT or PAT on a same security interface, then all traffic from the interface to a same security interface or an outside interface must match a NAT rule, as shown in Figure 42-4.

Figure 42-4 NAT Control and Same Security Traffic

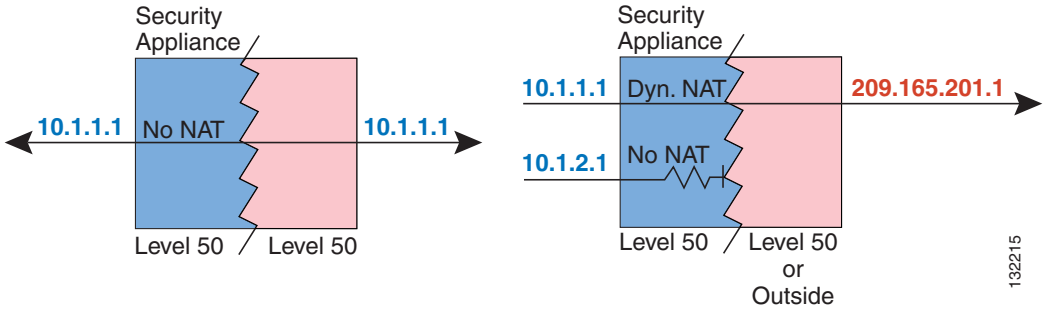

Similarly, if you enable outside dynamic NAT or PAT, then all outside traffic must match a NAT rule when it accesses an inside interface (see Figure 42-5).

Figure 42-5 NAT Control and Inbound Traffic

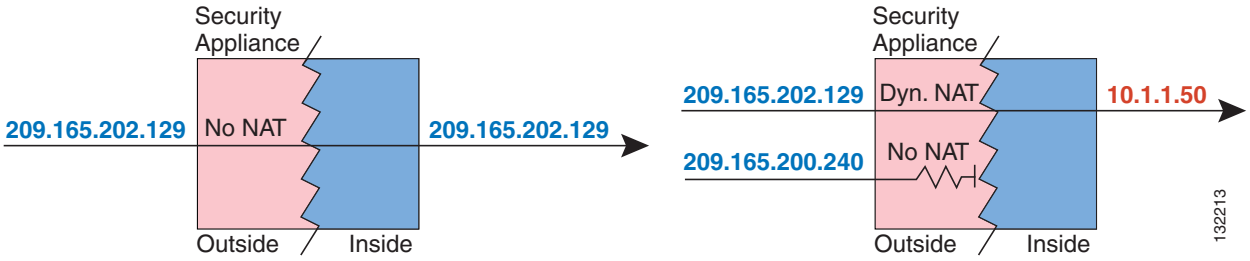

Static NAT does not cause these restrictions.

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 42-16 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 42-32 for more information).

To configure NAT control, see the "Configuring NAT Control" section on page 42-15.

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 ASA Classifies Packets" section on page 9-3 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 42-6
- PAT, page 42-8
- Static NAT, page 42-8
- Static PAT, page 42-9
- Bypassing NAT When NAT Control is Enabled, page 42-10

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 ASA 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 ASA 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 ASA. This condition occurs with an outbound access list, a management-only interface, or a backup interface in which the translation times out normally.

Figure 42-6 shows a remote host attempting to connect to the real address. The connection is denied, because the ASA only allows returning connections to the mapped address.

Figure 42-6 Remote Host Attempts to Connect to the Real Address

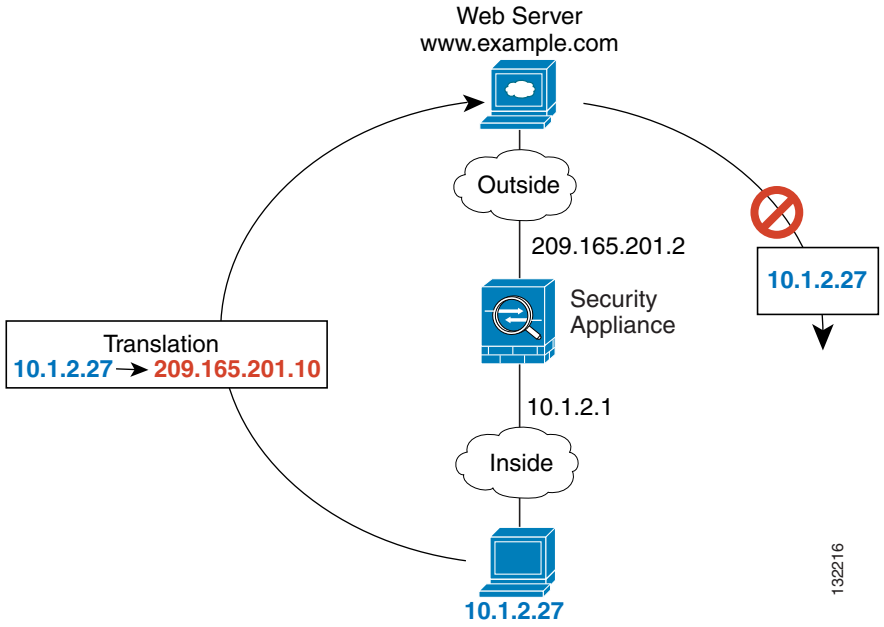

Figure 42-7 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the ASA drops the packet.

Figure 42-7 Remote Host Attempts to Initiate a Connection to a Mapped Address

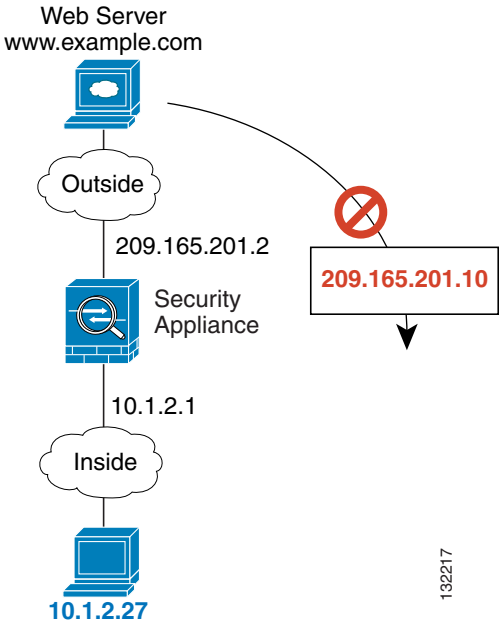

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access list allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case, you can rely on the security of the access list.

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 57-2 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 ASA 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 ASA 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 57-2 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 ASA 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 42-8).

Figure 42-8 Static PAT

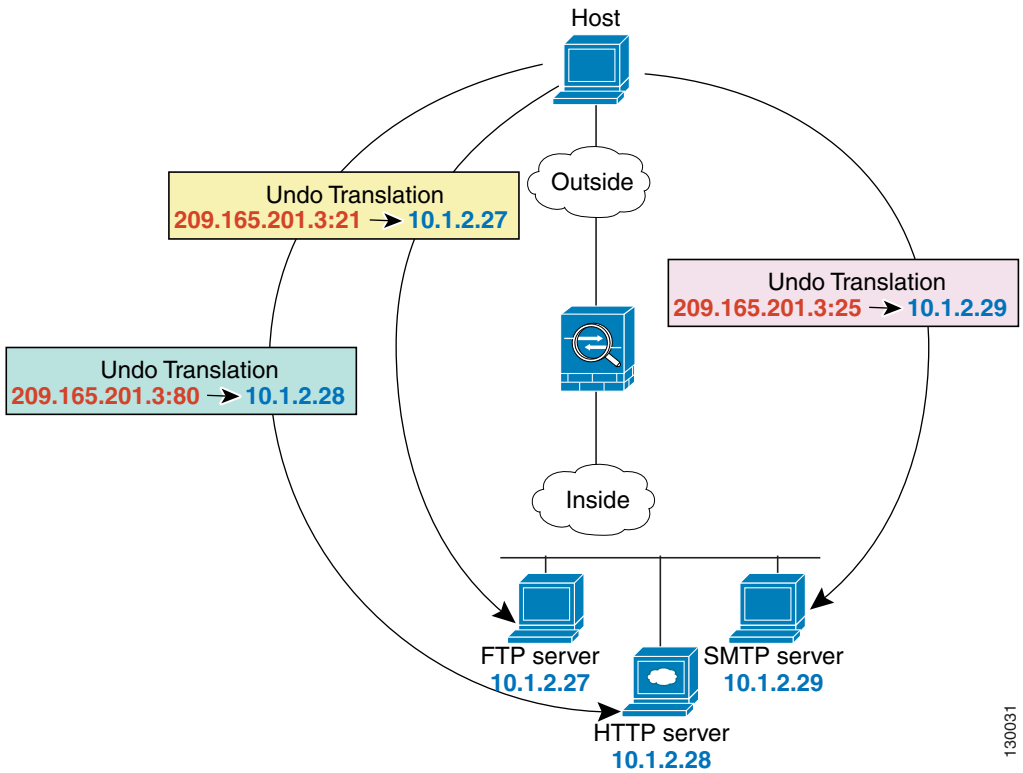

You can also use static PAT to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell 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 57-2 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 42-10 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 42-9 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 42-9 Policy NAT with Different Destination Addresses

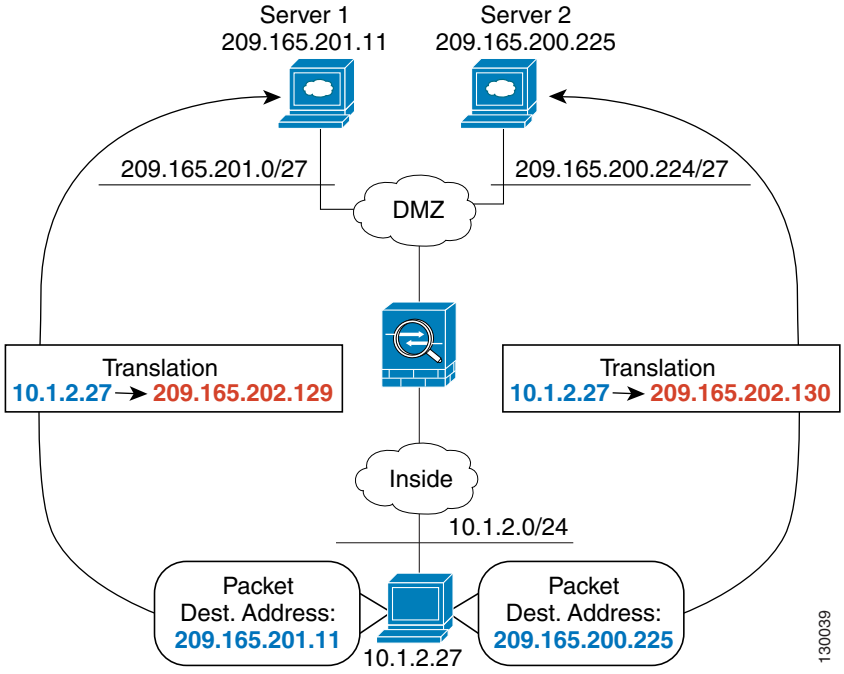

Figure 42-10 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 42-10 Policy NAT with Different Destination Ports

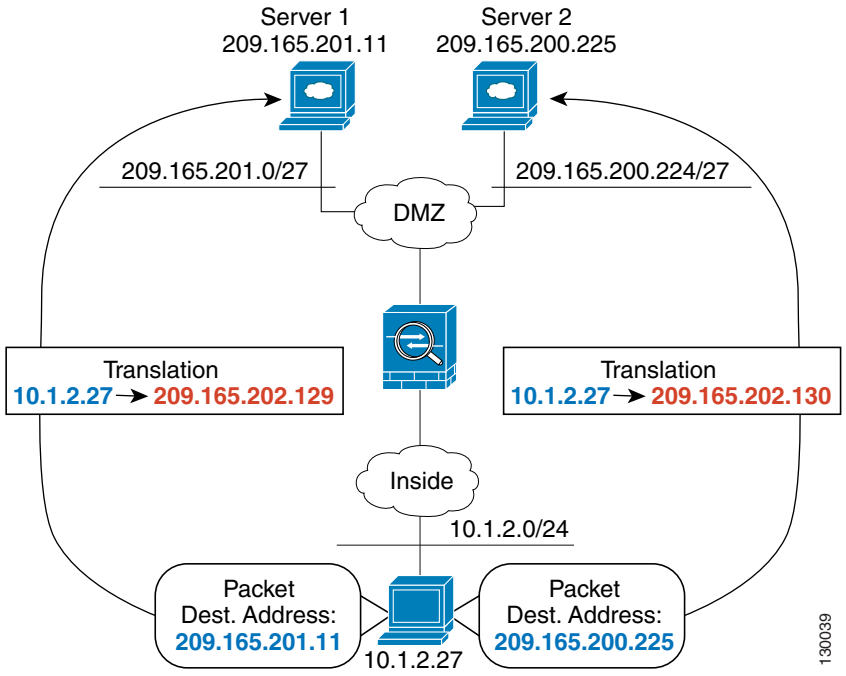

For policy static NAT, both translated and remote hosts can originate traffic. For traffic originated on the translated network, the NAT rule specifies the real addresses and the *destination* addresses, but for traffic originated on the remote network, the rule identifies the real addresses and the *source* addresses of remote hosts who are allowed to connect to the host using this translation.

Figure 42-11 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 42-11 Policy Static NAT with Destination Address Translation

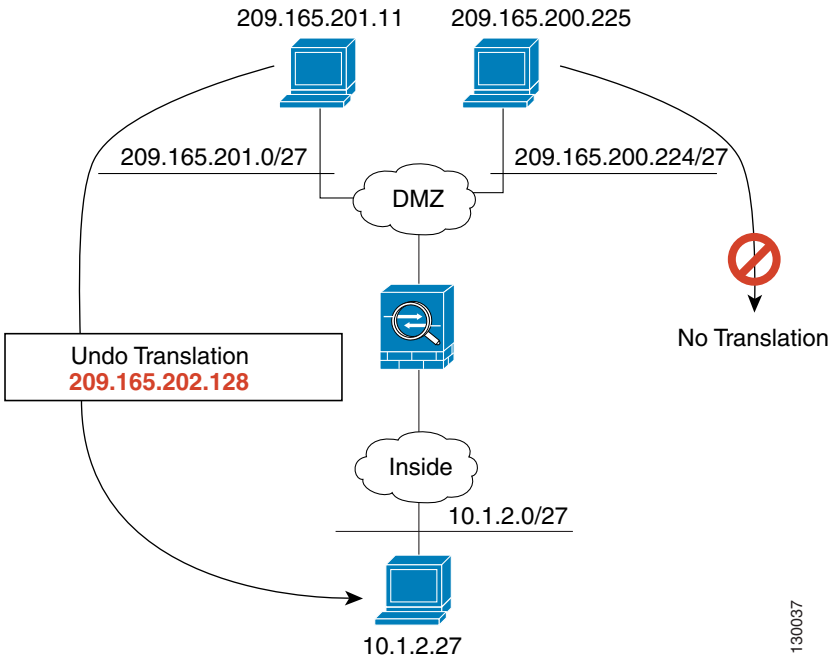

Note Policy NAT does not support SQL*Net, but it is supported by regular NAT. See the "When to Use Application Protocol Inspection" section on page 57-2 for information about NAT support for other protocols.

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 42-4 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 ASA 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 57-2 for supported inspection engines.

Order of NAT Rules Used to Match Real Addresses

The ASA 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 ASA.

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 ASA), the ASA 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 ASA 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 ASA 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 ASA 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 ASA.

DNS and NAT

You might need to configure the ASA to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation.

For example, a DNS server is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network (see Figure 42-12). 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 ASA 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 42-12 DNS Reply Modification

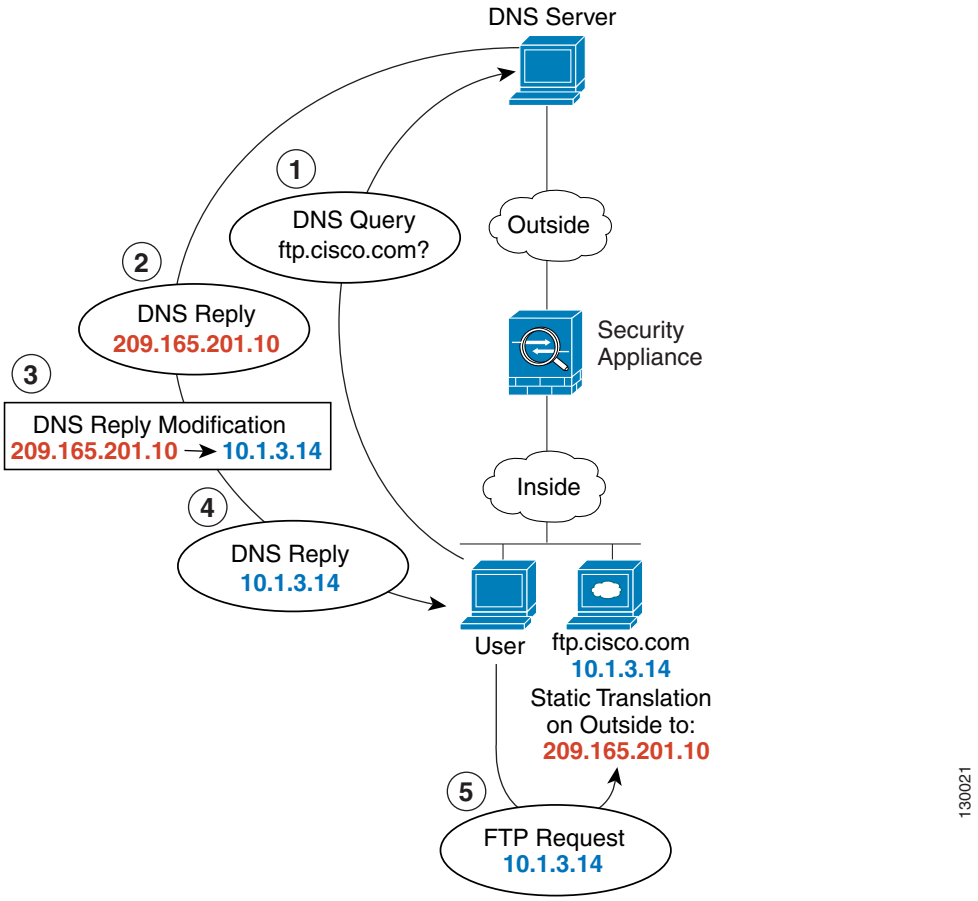

Note If a user on a different network (for example, DMZ) also requests the IP address for ftp.cisco.com from the outside DNS server, then the IP address in the DNS reply is also modified for this user, even though the user is not on the Inside interface referenced by the static rule.

Figure 42-13 shows a web server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 42-13 DNS Reply Modification Using Outside NAT

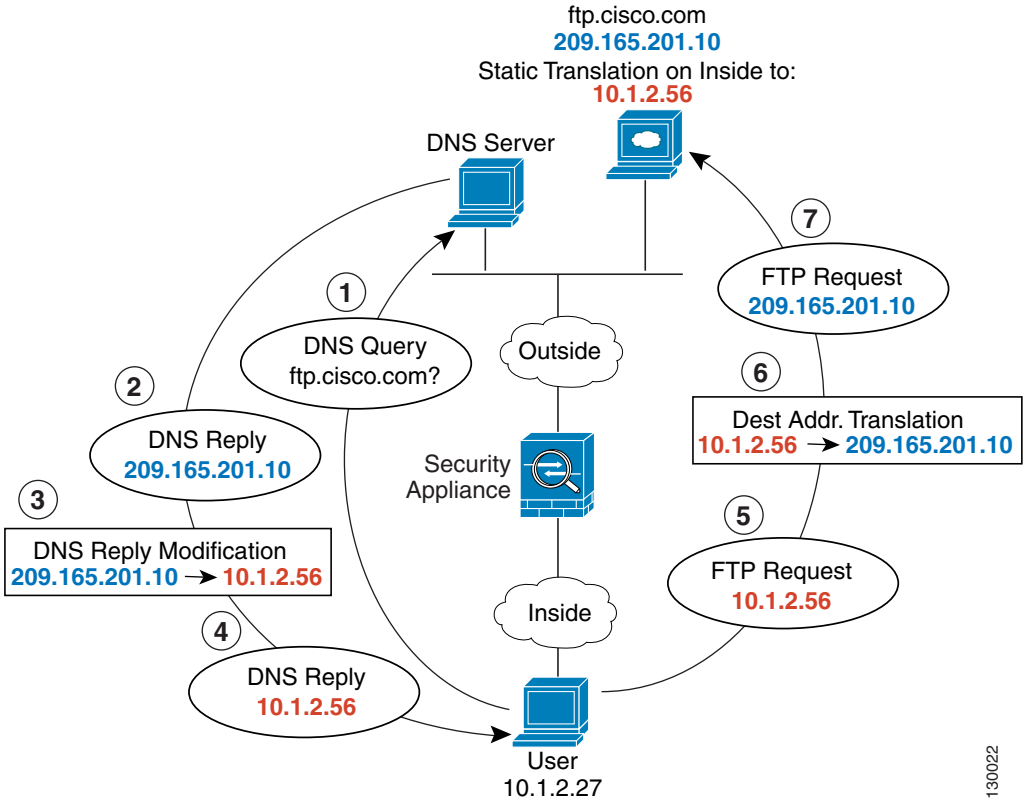

Configuring NAT Control

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 42-4 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 42-10 for more information.

This section includes the following topics:

- Dynamic NAT Implementation, page 42-16
- Managing Global Pools, page 42-21
- Configuring Dynamic NAT, PAT, or Identity NAT, page 42-22
- Configuring Dynamic Policy NAT or PAT, page 42-24

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 42-17
- NAT Rules on Different Interfaces with the Same Global Pools, page 42-17
- Global Pools on Different Interfaces with the Same Pool ID, page 42-18
- Multiple NAT Rules with Different Global Pools on the Same Interface, page 42-18
- Multiple Addresses in the Same Global Pool, page 42-19
- Outside NAT, page 42-20
- Real Addresses in a NAT Rule Must be Translated on All Lower or Same Security Interfaces, page 42-21

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 42-14).

Figure 42-14 NAT Rules on Multiple Interfaces Using the Same Global Pool

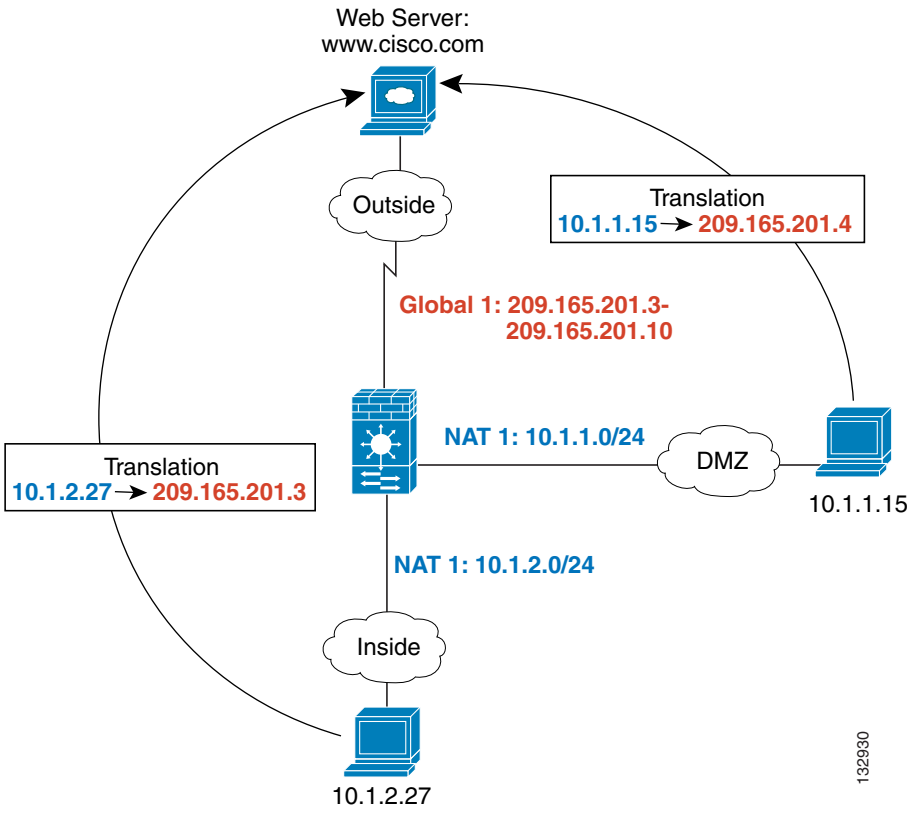

Global Pools on Different Interfaces with the Same Pool ID

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 42-15).

Figure 42-15 NAT Rules and Global Pools using the Same ID on Multiple Interfaces

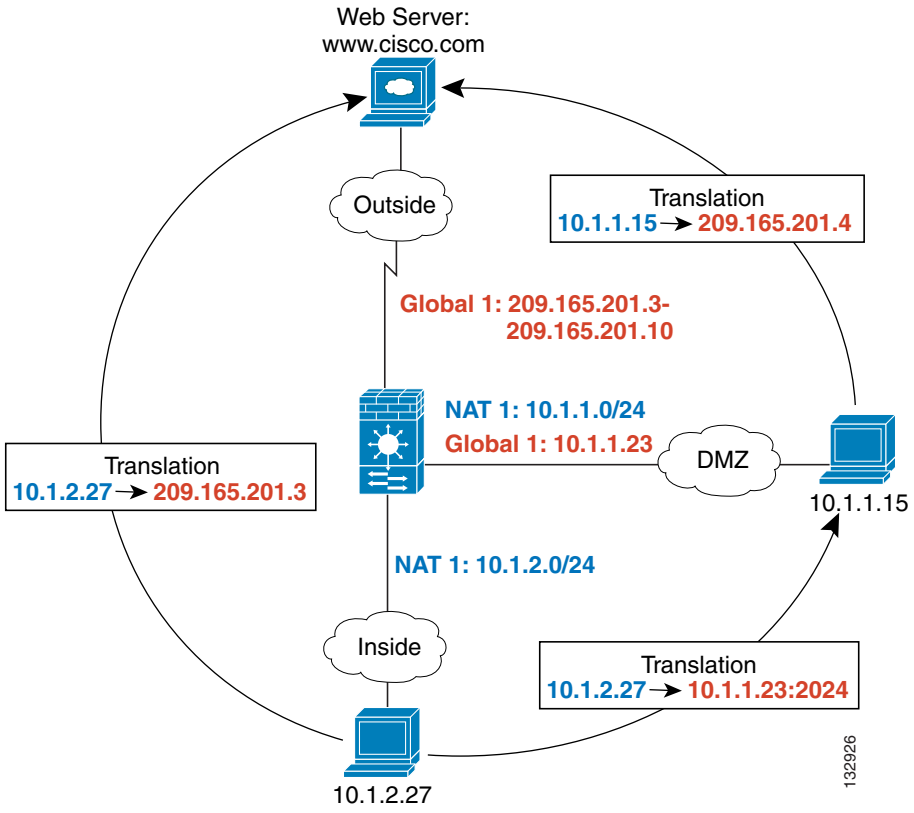

Multiple NAT Rules with Different Global Pools on the Same Interface

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 42-16). 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 42-16 Different NAT IDs

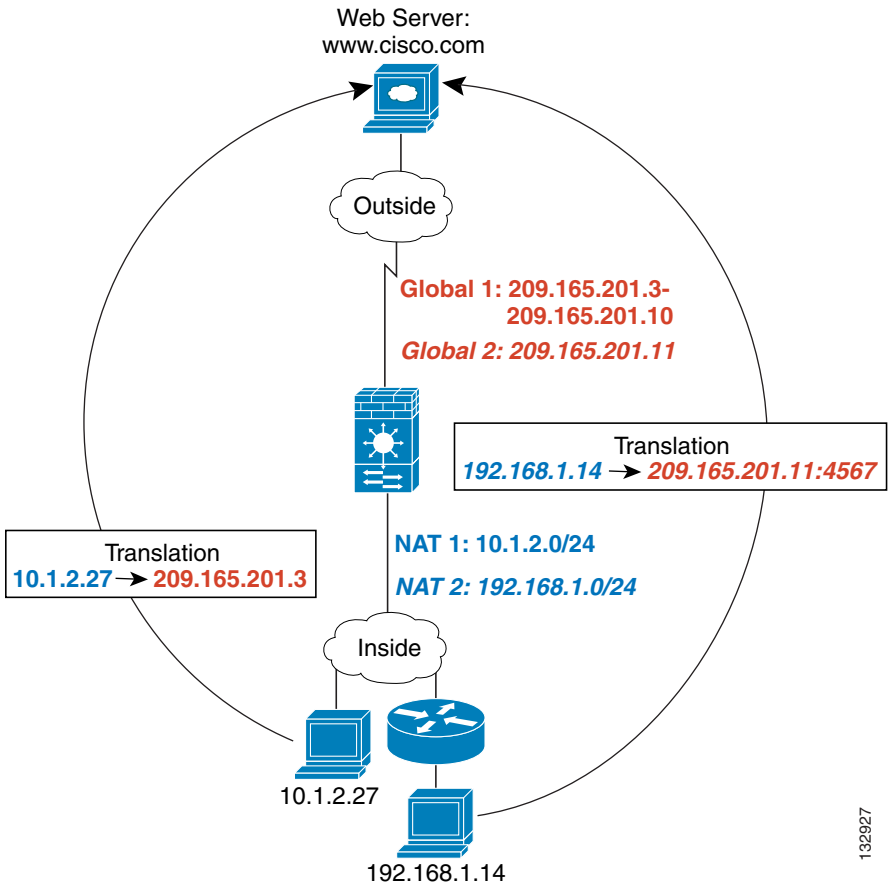

Multiple Addresses in the Same Global Pool

You can have multiple addresses in the same global pool; the ASA 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 42-17).

Figure 42-17 NAT and PAT Together

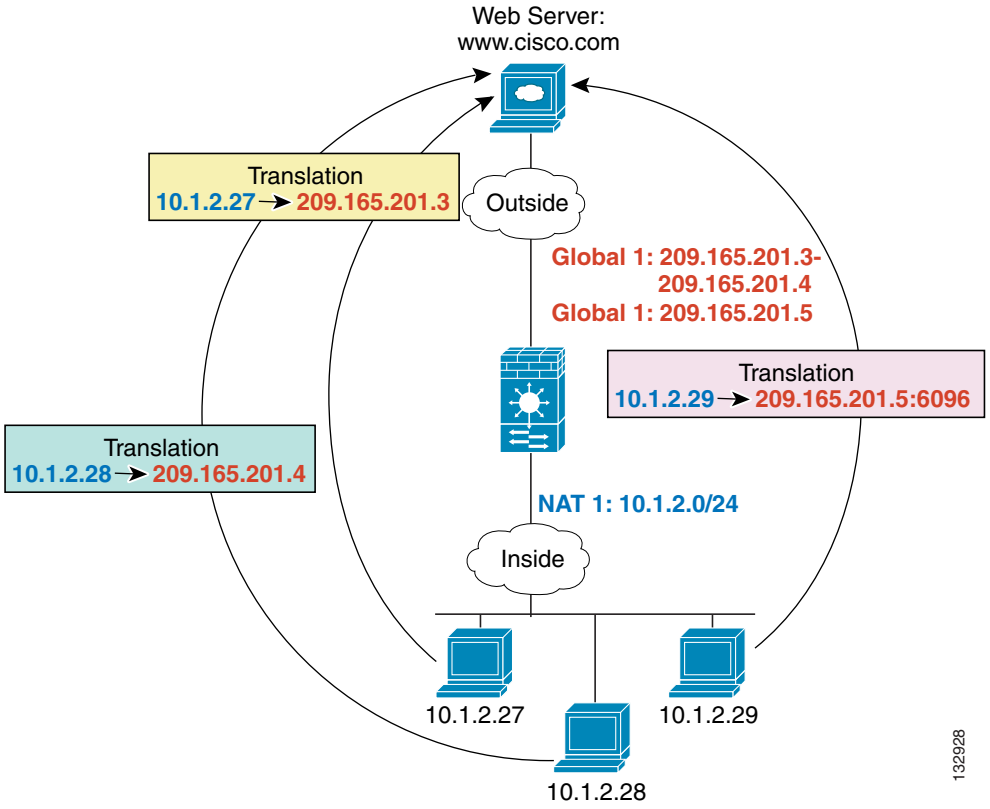

Outside NAT

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 42-18), 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 42-18 Outside NAT and Inside NAT Combined

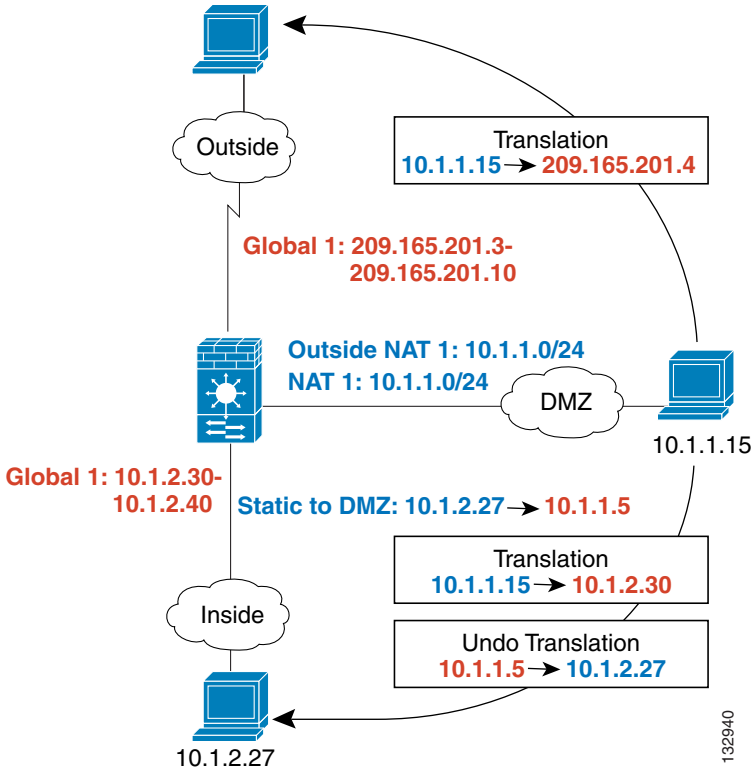

Real Addresses in a NAT Rule Must be Translated on All Lower or Same Security Interfaces

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 42-16.

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 ASA 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 42-19 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

Figure 42-19 shows typical dynamic NAT, dynamic PAT, and identity NAT scenarios. Only real hosts can initiate connections.

Figure 42-19 Dynamic NAT Scenarios

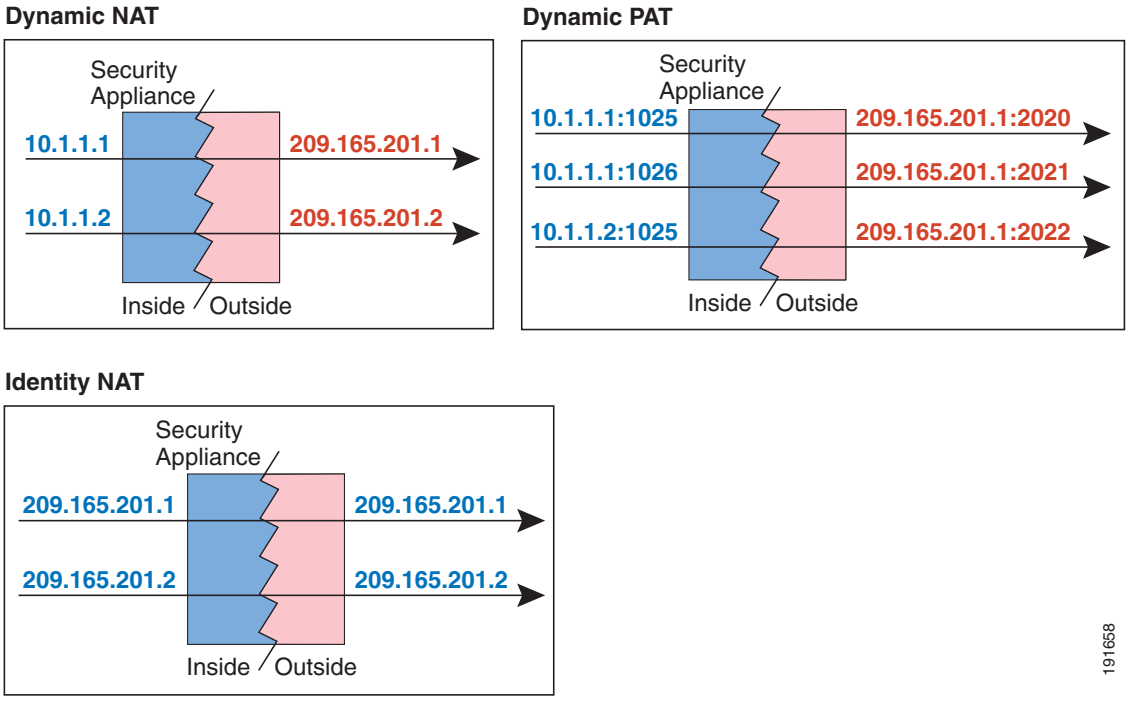

To configure a dynamic NAT, PAT, or identity NAT rule, perform the following steps.

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 ASA performs dynamic NAT. If the pool includes a single address, then the ASA 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 42-19 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 42-16.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 42-21.

- 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 42-13 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 Chapter 71, "Configuring Connection Settings"). If you set them in both places, then the ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA 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 ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

- **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

Figure 42-20 shows typical dynamic policy NAT and PAT scenarios. Only real hosts can initiate connections.

Figure 42-20 Dynamic Policy NAT Scenarios

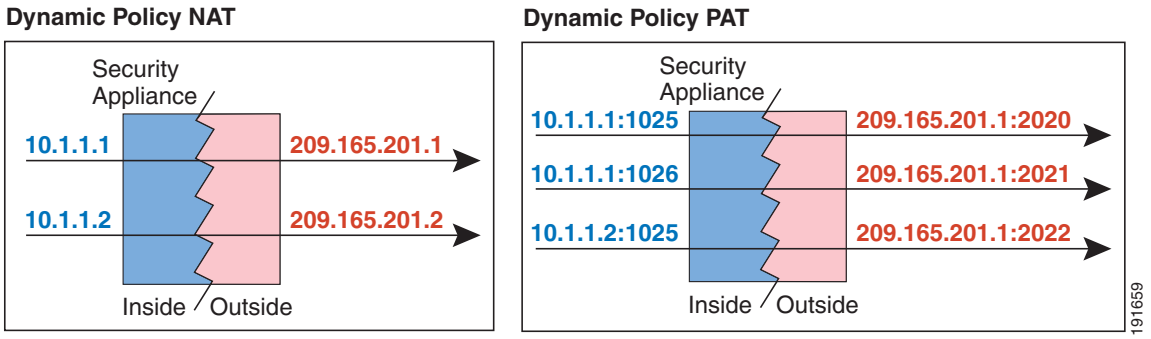

To configure dynamic policy NAT or PAT, perform the following steps:

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 ASA performs dynamic NAT. If the pool includes a single address, then the ASA 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 42-19 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 42-16.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 42-21.

- 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 42-13 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 ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA 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 ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

- **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 42-8.

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 42-10 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 ASA automatically translates the secondary ports. For more information about static PAT, see the "Static PAT" section on page 42-9.

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 42-27
- Configuring Static Policy NAT, PAT, or Identity NAT, page 42-30

Configuring Static NAT, PAT, or Identity NAT

Figure 42-21 shows typical static NAT, static PAT, and static identity NAT scenarios. The translation is always active so both translated and remote hosts can originate connections.

Figure 42-21 Static NAT Scenarios

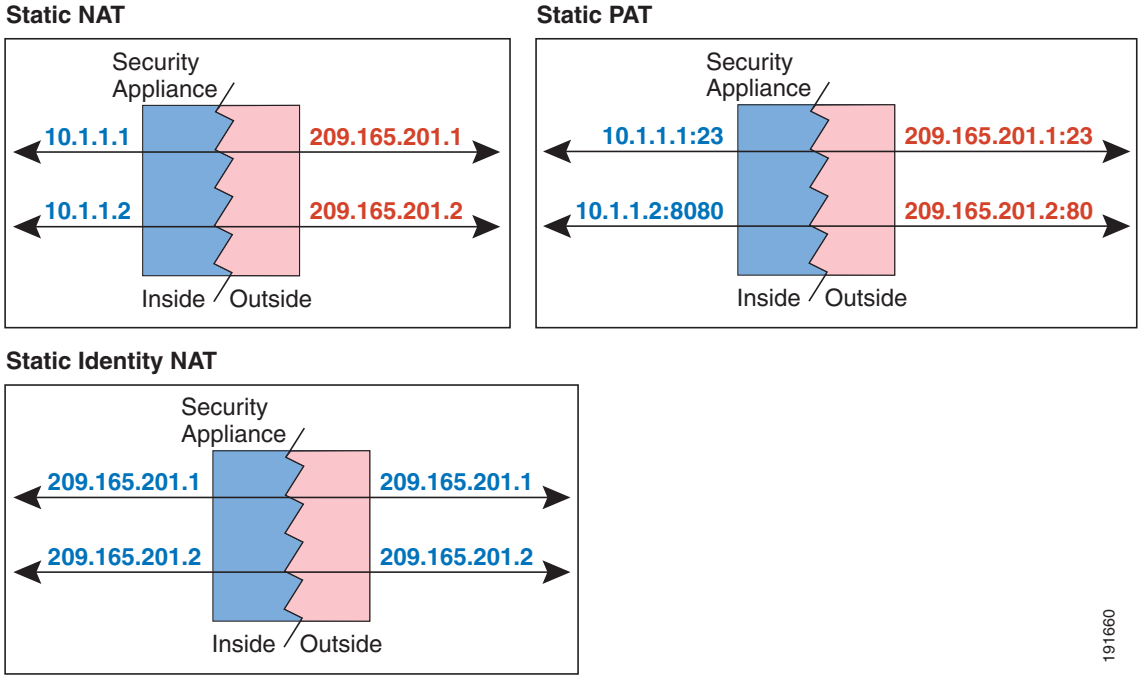

To configure static NAT, PAT, or identity NAT, perform the following steps:

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 42-13 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 ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA 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 ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.
- **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

Figure 42-22 shows typical static policy NAT, static policy PAT, and static policy identity NAT scenarios. The translation is always active so both translated and remote hosts can originate connections.

Figure 42-22 Static Policy NAT Scenarios

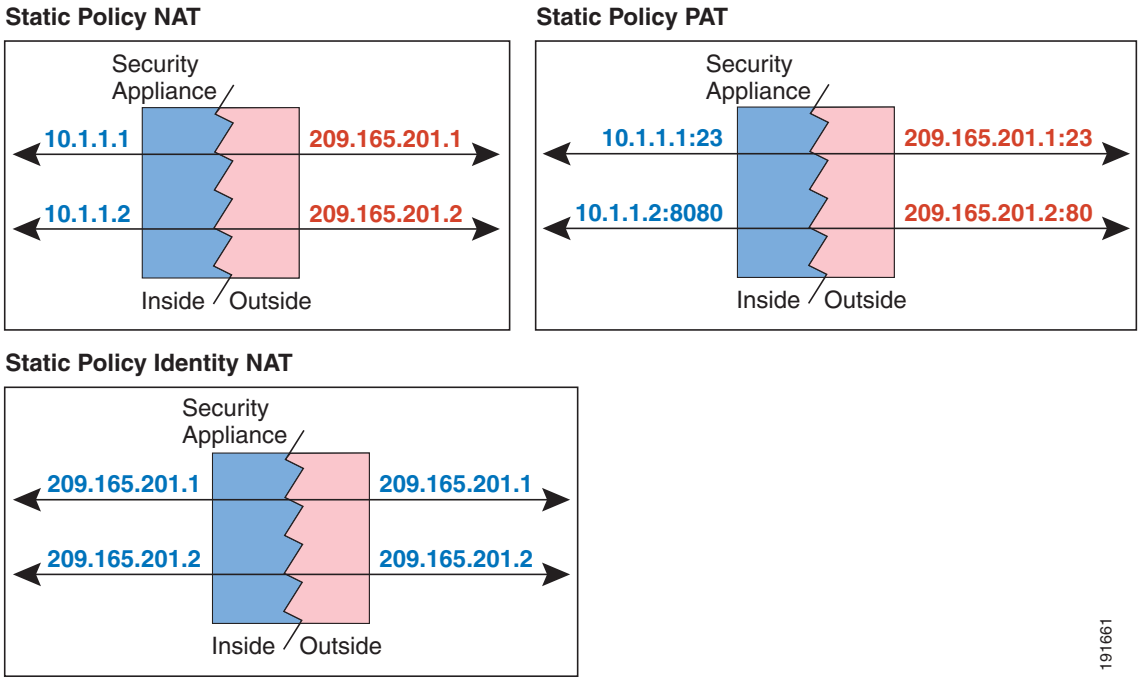

To configure static policy NAT, PAT, or identity NAT, perform the following steps:

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 42-13 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 ASA uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the ASA disables TCP sequence randomization.

- **Randomize sequence number**—With this check box checked (the default), the ASA 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 ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  – If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.

  – If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.

  - You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

- **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 42-10.

Figure 42-23 shows a typical NAT exemption scenario.

Figure 42-23 NAT Exemption

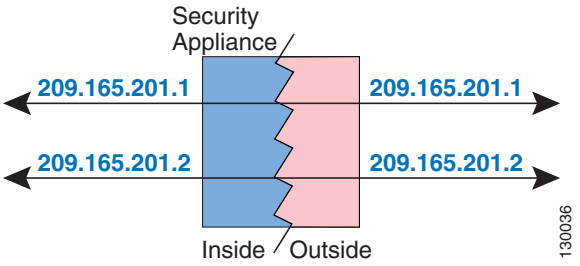

To configure NAT exemption, perform the following steps:

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 ASA decides whether to exempt a packet, the ASA 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.

PART 7

Configuring Service Policies

Configuring a Service Policy

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

This chapter includes the following sections:

- Information About Service Policies, page 44-1
- Licensing Requirements for Service Policies, page 44-5
- Guidelines and Limitations, page 44-6
- Default Settings, page 44-7
- Task Flows for Configuring Service Policies, page 44-8
- Adding a Service Policy Rule for Through Traffic, page 44-8
- Adding a Service Policy Rule for Management Traffic, page 44-13
- Managing the Order of Service Policy Rules, page 44-16
- Feature History for Service Policies, page 44-17

Information About Service Policies

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

- Supported Features, page 44-1
- Feature Directionality, page 44-2
- Feature Matching Within a Service Policy, page 44-3
- Order in Which Multiple Feature Actions are Applied, page 44-4
- Incompatibility of Certain Feature Actions, page 44-4
- Feature Matching for Multiple Service Policies, page 44-5

Supported Features

Table 44-1 lists the features supported by service policy rules.

Table 44-1 Service Policy Rule Features

Feature	For Through Traffic?	For Management Traffic?	See:
Application inspection (multiple types)	**All** except RADIUS accounting	**RADIUS accounting** only	• Chapter 57, "Getting Started with Application Layer Protocol Inspection." • Chapter 58, "Configuring Inspection of Basic Internet Protocols." • Chapter 59, "Configuring Inspection for Voice and Video Protocols." • Chapter 60, "Configuring Inspection of Database and Directory Protocols." • Chapter 61, "Configuring Inspection for Management Application Protocols." • Chapter 75, "Configuring the ASA for Cisco Cloud Web Security."
ASA CSC	**Yes**	No	Chapter 83, "Configuring the ASA CSC Module."
ASA IPS	**Yes**	No	Chapter 81, "Configuring the ASA IPS Module."
ASA CX	**Yes**	No	Chapter 82, "Configuring the ASA CX Module."
NetFlow Secure Event Logging filtering	**Yes**	**Yes**	Chapter 97, "Configuring NetFlow Secure Event Logging (NSEL)."
QoS input and output policing	**Yes**	No	Chapter 72, "Configuring QoS."
QoS standard priority queue	**Yes**	No	Chapter 72, "Configuring QoS."
QoS traffic shaping, hierarchical priority queue	**Yes**	**Yes**	Chapter 72, "Configuring QoS."
TCP and UDP connection limits and timeouts, and TCP sequence number randomization	**Yes**	**Yes**	Chapter 71, "Configuring Connection Settings."
TCP normalization	**Yes**	No	Chapter 71, "Configuring Connection Settings."
TCP state bypass	**Yes**	No	Chapter 71, "Configuring Connection Settings."
User statistics for Identity Firewall	**Yes**	**Yes**	See the **user-statistics** command in the command reference.

Feature Directionality

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

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

For features that are applied unidirectionally, for example QoS priority queue, only traffic that enters (or exits, depending on the feature) the interface to which you apply the policy map is affected. See Table 44-2 for the directionality of each feature.

Table 44-2 Feature Directionality

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

Feature Matching Within a Service Policy

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

1. A packet can match only one rule for an interface for each feature type.
2. When the packet matches a rule for a feature type, the ASA does not attempt to match it to any subsequent rules for that feature type.
3. If the packet matches a subsequent rule for a different feature type, however, then the ASA also applies the actions for the subsequent rule, if supported. See the "Incompatibility of Certain Feature Actions" section on page 44-4 for more information about unsupported combinations.

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

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

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

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

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

Order in Which Multiple Feature Actions are Applied

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

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

Actions are performed in the following order:

1. QoS input policing
2. TCP normalization, TCP and UDP connection limits and timeouts, TCP sequence number randomization, and TCP state bypass.

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

3. ASA CSC
4. Application inspections that can be combined with other inspections:
   a. IPv6
   b. IP options
   c. WAAS
5. Application inspections that cannot be combined with other inspections. See the "Incompatibility of Certain Feature Actions" section on page 44-4 for more information.
6. ASA IPS
7. ASA CX
8. QoS output policing
9. QoS standard priority queue
10. QoS traffic shaping, hierarchical priority queue

Incompatibility of Certain Feature Actions

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

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

- Most inspections should not be combined with another inspection, so the ASA only applies one inspection if you configure multiple inspections for the same traffic. HTTP inspection can be combined with the Cloud Web Security inspection. Other exceptions are listed in the "Order in Which Multiple Feature Actions are Applied" section on page 44-4.
- You cannot configure traffic to be sent to multiple modules, such as the ASA CX and ASA IPS.
- HTTP inspection is not compatible with the ASA CX.
- The ASA CX is not compatible with Cloud Web Security.

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

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

Feature Matching for Multiple Service Policies

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

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

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

Licensing Requirements for Service Policies

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

Traffic Class Guidelines

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

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

This limit also includes default traffic classes of all types, limiting user-configured traffic classes to approximately 235. See the "Default Traffic Classes" section on page 44-8.

Service Policy Guidelines

- Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP normalization, then both FTP inspection and TCP normalization are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.
- You can only apply one global policy. For example, you cannot create a global policy that includes feature set 1, and a separate global policy that includes feature set 2. All features must be included in a single policy.
- When you make service policy changes to the configuration, all *new* connections use the new service policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. **show** command output will not include data about the old connections.

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

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

Default Settings

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

- Default Configuration, page 44-7
- Default Traffic Classes, page 44-8

Default Configuration

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

The default policy includes the following application inspections:

- DNS 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
- IP Options

Default Traffic Classes

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

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

Task Flows for Configuring Service Policies

This section includes the following topics:

- Task Flow for Configuring a Service Policy Rule, page 44-8

Task Flow for Configuring a Service Policy Rule

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

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

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

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

Adding a Service Policy Rule for Through Traffic

See the "Supported Features" section on page 44-1 for more information. To add a service policy rule for through traffic, perform the following steps:

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

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

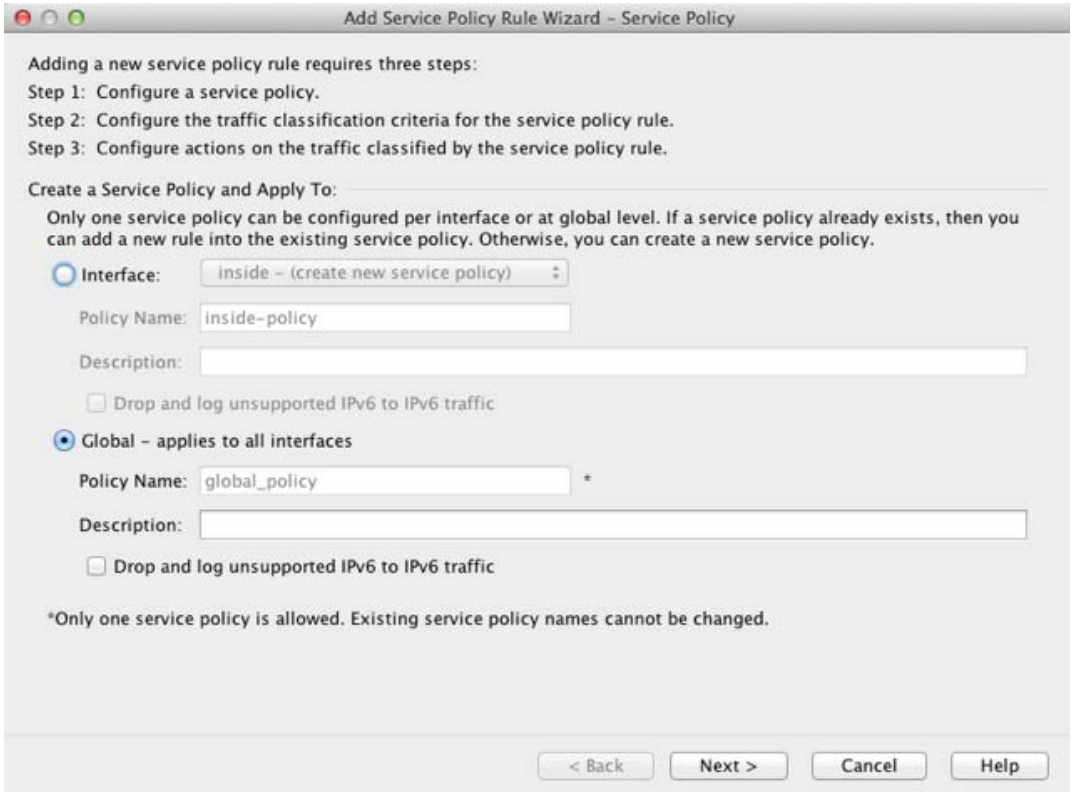

Note When you click the Add button, and not the small arrow on the right of the Add button, you add a through traffic rule by default. If you click the arrow on the Add button, you can choose between a through traffic rule and a management traffic rule.

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

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

     If you choose an interface that already has a policy, then the wizard lets you add a new service policy rule to the interface.
  b. If it is a new service policy, enter a name in the Policy Name field.
  c. (Optional) Enter a description in the Description field.
  d. (Optional) Check the **Drop and log unsupported IPv6 to IPv6 traffic** check box to generate a syslog (767001) for IPv6 traffic that is dropped by application inspections that do not support IPv6 traffic. By default, syslogs are not generated. For a list of inspections that support IPv6, see the "IPv6 Guidelines" section on page 44-6.

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See the "Default Settings" section on page 44-7 for more information. You can add a rule to the global policy using the wizard.

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

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

  c. (Optional) Check the **Drop and log unsupported IPv6 to IPv6 traffic** check box to generate a syslog (767001) for IPv6 traffic that is dropped by application inspections that do not support IPv6 traffic. By default, syslogs are not generated. For a list of inspections that support IPv6, see the "IPv6 Guidelines" section on page 44-6.

Step 3 Click **Next**.

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

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

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

  Identify the traffic using one of several criteria:

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

    This option, which is used in the default global policy, is a special shortcut that when used in a rule, ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same rule (See the "Incompatibility of Certain Feature Actions" section on page 44-4 for more information about combining actions). Normally, the ASA does not use the port number to determine the inspection applied, thus giving you the flexibility to apply inspections to non-standard ports, for example.

    See the "Default Settings" section on page 57-4 for a list of default ports. The ASA includes a default global policy that matches the default inspection traffic, and applies common inspections to the traffic on all interfaces. Not all applications whose ports are included in the Default Inspection Traffic class are enabled by default in the policy map.

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

  – **Source and Destination IP Address (uses ACL)**—The class matches traffic specified by an extended access list. If the ASA 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 44-16 for information about changing the order of ACEs.
- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).
- **Use class default as the traffic class**. This option uses the class-default class, which matches all traffic. The class-default class is created automatically by the ASA and placed at the end of the policy. If you do not apply any actions to it, it is still created by the ASA, but for internal purposes only. You can apply actions to this class, if desired, which might be more convenient than creating a new traffic class that matches all traffic. You can only create one rule for this service policy using the class-default class, because each traffic class can only be associated with a single rule per service policy.

Step 5 Click **Next**.

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

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

- Default Inspections—This dialog box is informational only, and shows the applications and the ports that are included in the traffic class.
- Source and Destination Address—This dialog box lets you set the source and destination addresses:

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

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

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

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

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

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

   Separate multiple addresses by a comma.

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

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

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

   Separate multiple addresses by a comma.

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

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

   By default, the service is IP.

   Separate multiple services by a comma.

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

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

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

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

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

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

   To add a new time range, click the **...** button. See the "Configuring Time Ranges" section on page 25-15 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 "Add or Edit an IPsec Remote Access Connection Profile" section on page 87-74 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 "Supported Features" section on page 44-1 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 ASA for management purposes. See the "Supported Features" section on page 44-1 for more information. This section includes the following topics:

Configuring a Service Policy Rule for Management Traffic

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

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

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

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

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

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

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

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

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

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

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See the "Default Settings" section on page 44-7 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 ASA 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 44-16 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 25-15 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 "Feature Directionality" section on page 44-2 for more information.

Step 10 To configure connection settings, see the "Configuring Connection Settings" section on page 71-8.

Step 11 Click **Finish**.

Managing the Order of Service Policy Rules

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

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

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

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

If your rule includes an access list with multiple ACEs, then the order of ACEs also affects the packet flow. The ASA tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an 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 44-1).

Figure 44-1 Moving an ACE

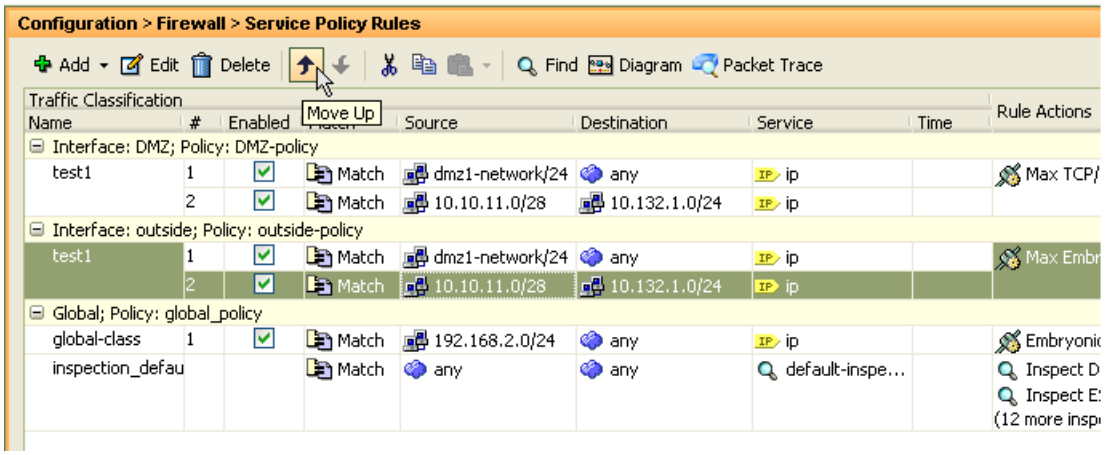

Note If you rearrange ACEs in an access list that is used in multiple service policies, then the change is inherited in all service policies.

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

Feature History for Service Policies

Table 44-3 lists the release history for this feature.

Table 44-3 Feature History for Service Policies

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

PART 8

Configuring AAA Servers and the Local Database

Configuring AAA Servers and the Local Database

This chapter describes support for authentication, authorization, and accounting (AAA, pronounced "triple A"), and how to configure AAA servers and the local database.

The chapter includes the following sections:

- Information About AAA, page 46-1
- Licensing Requirements for AAA Servers, page 46-9
- Guidelines and Limitations, page 46-10
- Configuring AAA, page 46-10
- Testing Server Authentication and Authorization, page 46-27
- Monitoring AAA Servers, page 46-27
- Additional References, page 46-28
- Feature History for AAA Servers, page 46-28

Information About AAA

AAA enables the ASA 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 connect through the ASA. (The Telnet server enforces authentication, too; the ASA 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:

- Information About Authentication, page 46-2
- Information About Authorization, page 46-2
- Information About Accounting, page 46-3
- Summary of Server Support, page 46-3
- RADIUS Server Support, page 46-4

- TACACS+ Server Support, page 46-5
- RSA/SDI Server Support, page 46-5
- NT Server Support, page 46-6
- Kerberos Server Support, page 46-6
- LDAP Server Support, page 46-6
- Local Database Support, Including as a Fallback Method, page 46-7
- How Fallback Works with Multiple Servers in a Group, page 46-8
- Using Certificates and User Login Credentials, page 46-8
- Task Flow for Configuring AAA, page 46-10

Information About Authentication

Authentication controls access by requiring valid user credentials, which are usually a username and password. You can configure the ASA to authenticate the following items:

- All administrative connections to the ASA, including the following sessions:
  - Telnet
  - SSH
  - Serial console
  - ASDM using HTTPS
  - VPN management access
- The **enable** command
- Network access
- VPN access

Information About Authorization

Authorization controls access *per user* after users are authenticated. You can configure the ASA to authorize the following items:

- Management commands
- Network access
- VPN access

Authorization controls the services and commands that are available to each authenticated user. If you did not 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 can authenticate inside users who try to access any server on the outside network and then limit the outside servers that a particular user can access using authorization.

The ASA caches the first 16 authorization requests per user, so if the user accesses the same services during the current authentication session, the ASA does not resend the request to the authorization server.

Information About Accounting

Accounting tracks traffic that passes through the ASA, 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 session start and stop times, username, the number of bytes that pass through the ASA for the session, the service used, and the duration of each session.

Summary of Server Support

Table 46-1 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, see the topics following the table.

Table 46-1 Summary of AAA Support

	Database Type							
AAA Service	**Local**	**RADIUS**	**TACACS+**	**SDI (RSA)**	**NT**	**Kerberos**	**LDAP**	**HTTP Form**
Authentication of...								
VPN users[1]	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes[2]
Firewall sessions	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Administrators	Yes	Yes	Yes	Yes[3]	Yes	Yes	Yes	No
Authorization of...								
VPN users	Yes	Yes	No	No	No	No	Yes	No
Firewall sessions	No	Yes[4]	Yes	No	No	No	No	No
Administrators	Yes[5]	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[6]	Yes	No	No	No	No	No

1. For SSL VPN connections, either PAP or MS-CHAPv2 can be used.
2. HTTP Form protocol supports both authentication and SSO operations for clientless SSL VPN users sessions only.
3. RSA/SDI is supported for ASDM HTTP administrative access with ASA 5500 software version 8.2(1) or later.
4. For firewall sessions, RADIUS authorization is supported with user-specific access lists only, which are received or specified in a RADIUS authentication response.
5. Local command authorization is supported by privilege level only.
6. Command accounting is available for TACACS+ only.

Note In addition to the native protocol authentication listed in Table 46-1, the ASA supports proxying authentication. For example, the ASA can proxy to an RSA/SDI and/or LDAP server via a RADIUS server. Authentication via digital certificates and/or digital certificates with the AAA combinations listed in the table are also supported.

RADIUS Server Support

The ASA supports the following RFC-compliant RADIUS servers for AAA:

- Cisco Secure ACS 3.2, 4.0, 4.1, 4.2, and 5.x
- Cisco Identity Services Engine (ISE)
- RSA RADIUS in RSA Authentication Manager 5.2, 6.1, and 7.x
- Microsoft

Authentication Methods

The ASA supports the following authentication methods with RADIUS:

- PAP—For all connection types.
- CHAP and MS-CHAPv1—For L2TP-over-IPsec connections.
- MS-CHAPv2—For L2TP-over-IPsec connections, and for regular IPsec remote access connections when the password management feature is enabled. You can also use MS-CHAPv2 with clientless connections.
- Authentication Proxy modes—Including RADIUS to Active Directory, RADIUS to RSA/SDI, RADIUS to Token-server, and RSA/SDI to RADIUS connections,

Note To enable MS-CHAPv2 as the protocol used between the ASA and the RADIUS server for a VPN connection, password management must be enabled in the tunnel group general attributes. Enabling password management generates an MS-CHAPv2 authentication request from the ASA to the RADIUS server. See the description of the **password-management** command for details.

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 using the **no mschapv2-capable** command.

Attribute Support

The ASA 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 Vendor-Specific Attributes (VSAs), identified by RADIUS vendor ID 9.
- Cisco VPN-related VSAs, identified by RADIUS vendor ID 3076.
- Microsoft VSAs, defined in RFC 2548.
- Cisco VSA (Cisco-Priv-Level), which provides a standard 0-15 numeric ranking of privileges, with 1 being the lowest level and 15 being the highest level. A zero level indicates no privileges. The first level (login) allows privileged EXEC access for the commands available at this level. The second level (enable) allows CLI configuration privileges.

- A list of attributes is available at the following URL: http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/ref_extserver.html#wp1605508

RADIUS Authorization Functions

The ASA can use RADIUS servers for user authorization of VPN remote access and firewall cut-through-proxy sessions 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 ASA. Access to a given service is either permitted or denied by the access list. The ASA deletes the access list when the authentication session expires.

In addition to access lists, the ASA supports many other attributes for authorization and setting of permissions for VPN remote access and firewall cut-through proxy sessions. For a complete list of authorization attributes, see the following URL: http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/ref_extserver.html#wp1605508

TACACS+ Server Support

The ASA supports TACACS+ authentication with ASCII, PAP, CHAP, and MS-CHAPv1.

RSA/SDI Server Support

The RSA SecureID servers are also known as SDI servers.

This section includes the following topics:

- RSA/SDI Version Support, page 46-5
- Two-step Authentication Process, page 46-5
- RSA/SDI Primary and Replica Servers, page 46-6

RSA/SDI Version Support

The ASA supports SDI Versions 5.x, 6.x, and 7.x. 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 or Server IP address, with .sdi appended.

A version 5.x, 6.x, or 7.x SDI server that you configure on the ASA can be either the primary or any one of the replicas. See the "RSA/SDI Primary and Replica Servers" section on page 46-6 for information about how the SDI agent selects servers to authenticate users.

Two-step Authentication Process

SDI Versions 5.x, 6.x, or 7.x use 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 actions means that the same user cannot authenticate to two ASAs using the same authentication servers simultaneously. After a successful username lock, the ASA sends the passcode.

RSA/SDI Primary and Replica Servers

The ASA obtains the server list when the first user authenticates to the configured server, which can be either a primary or a replica. The ASA then assigns priorities to each of the servers on the list, and subsequent server selection is derived at random from those assigned priorities. The highest priority servers have a higher likelihood of being selected.

NT Server Support

The ASA 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, which is a limitation of NTLM Version 1.

Kerberos Server Support

The ASA supports 3DES, DES, and RC4 encryption types.

Note The ASA 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 ASA.

LDAP Server Support

The ASA supports LDAP. This section includes the following topics:

- Authentication with LDAP, page 46-6
- LDAP Server Types, page 46-7

Authentication with LDAP

During authentication, the ASA acts as a client proxy to the LDAP server for the user, and authenticates to the LDAP server in either plain text or by using the SASL protocol. By default, the ASA passes authentication parameters, usually a username and password, to the LDAP server in plain text.

The ASA supports the following SASL mechanisms, listed in order of increasing strength:

- Digest-MD5—The ASA responds to the LDAP server with an MD5 value computed from the username and password.
- Kerberos—The ASA responds to the LDAP server by sending the username and realm using the GSSAPI Kerberos mechanism.

You can configure the ASA and LDAP server to support any combination of these SASL mechanisms. If you configure multiple mechanisms, the ASA retrieves the list of SASL mechanisms that are configured on the server and sets the authentication mechanism to the strongest mechanism configured on both the ASA and the server. For example, if both the LDAP server and the ASA support both mechanisms, the ASA selects Kerberos, the stronger of the mechanisms.

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 that is applied to the VPN session. Thus, using LDAP accomplishes authentication and authorization in a single step.

LDAP Server Types

The ASA 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, Novell, OpenLDAP, and other LDAPv3 directory servers.

By default, the ASA auto-detects whether it is connected to Microsoft Active Directory, Sun LDAP, Novell, OpenLDAP, 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.

When configuring the server type, note the following guidelines:

- The DN configured on the ASA 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 ACL on the default password policy.
- You must configure LDAP over SSL to enable password management with Microsoft Active Directory and Sun servers.
- The ASA does not support password management with Novell, OpenLDAP, and other LDAPv3 directory servers.
- The ASA uses the Login Distinguished Name (DN) and Login Password to establish a trust relationship (bind) with an LDAP server. For more information, see the "Binding the ASA to the LDAP Server" section on page B-4.

HTTP Forms Authentication for Clientless SSL VPN

The ASA can use the HTTP Form protocol for both authentication and SSO operations of Clientless SSL VPN user sessions only.

Local Database Support, Including as a Fallback Method

The ASA maintains a local database that you can populate with user profiles.

The local database can act as a fallback method for several functions. This behavior is designed to help you prevent accidental lockout from the ASA.

For users who need fallback support, we recommend that their usernames and passwords in the local database match their usernames and passwords on the AAA servers. This practice 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—If the servers in the group are all unavailable, the ASA uses the local database to authenticate administrative access, which can also include enable password authentication.
- Command authorization—If the TACACS+ servers in the group are all 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 ASA if AAA servers that normally support these VPN services are unavailable. When a 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.

How Fallback Works with Multiple Servers in a Group

If you configure multiple servers in a server group and you enable fallback to the local database for the server group, fallback occurs when no server in the group responds to the authentication request from the ASA. To illustrate, consider this scenario:

You configure an LDAP server group with two Active Directory servers, server 1 and server 2, in that order. When the remote user logs in, the ASA attempts to authenticate to server 1.

If server 1 responds with an authentication failure (such as *user not found*), the ASA does not attempt to authenticate to server 2.

If server 1 does not respond within the timeout period (or the number of authentication attempts exceeds the configured maximum), the ASA tries server 2.

If both servers in the group do not respond, and the ASA is configured to fall back to the local database, the ASA tries to authenticate to the local database.

Using Certificates and User Login Credentials

The following section describes the different methods of using certificates and user login credentials (username and password) for authentication and authorization. These methods apply to IPsec, AnyConnect, and Clientless SSL VPN.

In all cases, LDAP authorization does not use the password as a credential. RADIUS authorization uses either a common password for all users or the username as a password.

This section includes the following topics:

- Using User Login Credentials, page 46-8
- Using Certificates, page 46-9

Using User Login Credentials

The default method for authentication and authorization uses the user login credentials.

- Authentication

  - Enabled by the authentication server group setting in the tunnel group (also called ASDM Connection Profile)
  - Uses the username and password as credentials
- Authorization
  - Enabled by the authorization server group setting in the tunnel group (also called ASDM Connection Profile)
  - Uses the username as a credential

Using Certificates

If user digital certificates are configured, the ASA first validates the certificate. It does not, however, use any of the DNs from certificates as a username for the authentication.

If both authentication and authorization are enabled, the ASA uses the user login credentials for both user authentication and authorization.

- Authentication
  - Enabled by the authentication server group setting
  - Uses the username and password as credentials
- Authorization
  - Enabled by the authorization server group setting
  - Uses the username as a credential

If authentication is disabled and authorization is enabled, the ASA uses the primary DN field for authorization.

- Authentication
  - DISABLED (set to None) by the authentication server group setting
  - No credentials used
- Authorization
  - Enabled by the authorization server group setting
  - Uses the username value of the certificate primary DN field as a credential

Note If the primary DN field is not present in the certificate, the ASA uses the secondary DN field value as the username for the authorization request.

For example, consider a user certificate that includes the following Subject DN fields and values:

```
Cn=anyuser,OU=sales;O=XYZCorporation;L=boston;S=mass;C=us;ea=anyuser@example.com
```

If the Primary DN = EA (E-mail Address) and the Secondary DN = CN (Common Name), then the username used in the authorization request would be anyuser@example.com.

Licensing Requirements for AAA Servers

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.

Configuring AAA

This section includes the following topics:

- Configuring AAA Server Groups, page 46-11
- Adding a Server to a Group, page 46-12
- Configuring AAA Server Parameters, page 46-13
- Configuring LDAP Attribute Maps, page 46-19
- Adding a User Account to the Local Database, page 46-20
- Authenticating Users with a Public Key for SSH, page 46-25
- Adding an Authentication Prompt, page 46-26

Task Flow for Configuring AAA

Step 1 Do one or both of the following:

- Add a AAA server group. See the "Configuring AAA Server Groups" section on page 46-11.
- Add a user to the local database. See the "Adding a User Account to the Local Database" section on page 46-20.

Step 2 For a server group, add a server to the group. See the "Adding a Server to a Group" section on page 46-12.

Step 3 For a server group, configure server parameters. See the "Configuring AAA Server Parameters" section on page 46-13.

Step 4 For an LDAP server, configure LDAP attribute maps. See the "Configuring LDAP Attribute Maps" section on page 46-19.

Step 5 (Optional) Specify text to display to the user during the AAA authentication challenge process. See the "Adding an Authentication Prompt" section on page 46-26.

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+.

Guidelines

- 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 ASA 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 ASA continues to try the AAA servers.

Detailed Steps

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, enter 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 ASA sends accounting data to only one server.

In Simultaneous mode, the ASA 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 the Depletion reactivation mode, enter a time interval in the Dead Time field.

The Dead Time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers.

Step 8 In the Max Failed Attempts field, add the number of failed attempts allowed.

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, perform the following steps:

a. Check the **Enable interim accounting update** check box if you want to enable multi-session accounting for clientless SSL and AnyConnect sessions.

b. Check the **Enable Active Directory Agent Mode** check box to specify the shared secret between the ASA and the AD agent and indicate that a RADIUS server group includes AD agents that are not full-function RADIUS servers. Only a RADIUS server group that has been configured using this option can be associated with user identity.

c. Click the **VPN3K Compatibility Option** down arrow to expand the list, and click one of the following radio buttons to specify whether or not a downloadable ACL received from a RADIUS packet 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 Add AAA Server Group dialog box closes, and the new 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 to the running configuration.

Adding a Server to a Group

To add a AAA server to a group, perform the following steps.

Detailed 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 list, choose the interface name on which 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 that 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 ASA 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 that are unique to each server type:

- RADIUS Server Fields, page 46-14
- TACACS+ Server Fields, page 46-15
- SDI Server Fields, page 46-15
- Windows NT Domain Server Fields, page 46-15
- Kerberos Server Fields, page 46-16
- LDAP Server Fields, page 46-16
- HTTP Form Server Fields, page 46-18

Step 7 Click **OK**.

The Add AAA Server Group 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 to the running configuration.

Configuring AAA Server Parameters

This section lists the unique fields for each server type when you add a server to a server group and includes the following topics:

- RADIUS Server Fields, page 46-14
- TACACS+ Server Fields, page 46-15
- SDI Server Fields, page 46-15
- Windows NT Domain Server Fields, page 46-15
- Kerberos Server Fields, page 46-16
- LDAP Server Fields, page 46-16
- HTTP Form Server Fields, page 46-18

For more information, see the "Adding a Server to a Group" section on page 46-12.

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 46-12.

Field	Description
ACL Netmask Convert	How you want the ASA to handle netmasks received in downloadable access lists. • Detect automatically: The ASA attempts to determine the type of netmask expression used. If the ASA detects a wildcard netmask expression, the ASA 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 ASA assumes downloadable access lists received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed. • Wildcard: The ASA 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.
Common Password	A case-sensitive password that is common among users who access this RADIUS authorization server through this ASA. 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 user 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.
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.
Retry Interval	The duration of time, 1 to 10 seconds, that the ASA waits between attempts to contact the server.
Server Accounting Port	The server port to be used for accounting of users. The default port is 1646.

Field	Description
Server Authentication Port	The server port to be used for authentication of users. The default port is 1645.
Server Secret Key	The shared secret key used to authenticate the RADIUS server to the ASA. The server secret that 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.

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 46-12.

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 ASA. The server secret that you configure here should match the one that is 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 46-12.

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 ASA 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 46-12.

Field	Description
Server Port	Port number 139, or the TCP port number used by the ASA 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 46-12.

Field	Description
Server Port	Server port number 88, or the UDP port number over which the ASA communicates with the Kerberos server.
Retry Interval	The duration of time, 1 to 10 seconds, that the ASA 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 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 46-12.

Field	Description
Enable LDAP over SSL check box	When checked, SSL secures communications between the ASA and the LDAP server. Also called secure LDAP (LDAP-S). **Note** If you do not configure the SASL protocol, we strongly recommend that you secure LDAP communications with SSL.
Server Port	TCP port number 389, the port which the ASA 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 a VPN remote access password management capability, which requires LDAP-S.

Field	Description
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, 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 (that is, searches 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).
Login DN	The ASA 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 ASA authenticates to the server using the Login DN and the Login password. When performing a Microsoft Active Directory read-only operation (such as authentication, authorization, or group-search), the ASA can bind with a Login DN with fewer 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. The following is an example of a Login DN: `cn=Binduser1,ou=Admins,ou=Users,dc=company_A,dc=com` The ASA 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 ASA does not support anonymous authentication.
Login Password	The password for the Login DN user account. The characters you type are replaced with asterisks.

Field	Description
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. For more information, see the "Adding an Authentication Prompt" section on page 46-26.
SASL MD5 authentication check box	When checked, the MD5 mechanism of the SASL authenticates communications between the ASA and the LDAP server.
SASL Kerberos authentication	When checked, the Kerberos mechanism of the SASL secures authentication communications between the ASA 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.
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 ASA uses the Base DN for AD group retrieval. ASDM uses the list of retrieved AD groups to define AAA selection criteria for dynamic access policies. For more information, see the **show ad-groups** command.
Group Search Timeout	Specifies the maximum time to wait for a response from an AD server that was 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 ASA.

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 46-12.

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.

Field	Description
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.
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.

Configuring LDAP Attribute Maps

The ASA can use an LDAP directory for authenticating VPN remote access users or firewall network access/cut-through-proxy sessions and/or for setting policy permissions (also called authorization attributes), such as ACLs, bookmark lists, DNS or WINS settings, session timers, and so on. That is, you can set the key attributes that exist in a local group policy externally through an LDAP server.

The authorization process is accomplished by means of LDAP attribute maps (similar to a RADIUS dictionary that defines vendor-specific attributes), which translate the native LDAP user attributes to Cisco ASA attribute names. You can then bind these attribute maps to LDAP servers or remove them, as needed. You can also show or clear attribute maps.

Guidelines

The ldap-attribute-map has a limitation with multi-valued attributes. For example, if a user is a memberOf of several AD groups and the ldap attribute map matches on more than one of them, the mapped value is chosen based on the alphabetization of the matched entries.

To use the attribute mapping features correctly, you need to understand Cisco LDAP attribute names and values, as well as the user-defined attribute names and values. For more information about LDAP attribute maps, see the "Active Directory/LDAP VPN Remote Access Authorization Examples" section on page B-15.

The names of frequently mapped Cisco LDAP attributes and the type of user-defined attributes that they would commonly be mapped to include the following:

- IETF-Radius-Class (Group_Policy in ASA version 8.2 and later)—Sets the group policy based on the directory's department or user group (for example, Microsoft Active Directory memberOf) attribute value. The group-policy attribute replaced the IETF-Radius-Class attribute with ASDM version 6.2/ASA version 8.2 or later.
- IETF-Radius-Filter-Id—An access control list or ACL applied to VPN clients, IPsec, and SSL.
- IETF-Radius-Framed-IP-Address—Assigns a static IP address assigned to a VPN remote access client, IPsec, and SSL.
- Banner1—Displays a text banner when the VPN remote access user logs in.
- Tunneling-Protocols—Allows or denies the VPN remote access session based on the access type.

Note A single ldap attribute map may contain one or many attributes. You can only assign one ldap attribute to a specific LDAP server.

To map LDAP features correctly, perform the following steps:

Detailed Steps

Step 1 Choose **Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map**, and then 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 the corresponding attribute of your organization.

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 to which the value specified in the previous step maps.

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 to the running configuration.

Adding a User Account to the Local Database

This section describes how to manage users in the local database.

This section includes the following topics:

- Adding a User, page 46-21
- Configuring VPN Policy Attributes for a User, page 46-23

Adding a User

To add a user to the local database, perform the following steps:

To add a user to the local database, perform the following steps:

Guidelines

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 the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1). 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.
- 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 ASA refers to the user privilege level to determine which 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

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.

Limitations

You cannot use the local database for network access authorization.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Users/AAA > User Accounts**, and then click **Add**.

The Add User Account-Identity dialog box appears.

Step 2 In the Username field, enter a username from 4 to 64 characters long.

Step 3 In the Password field, enter a password between 3 and 32 characters. Passwords are case-sensitive. The field displays only asterisks. To protect security, we recommend a password length of at least 8 characters.

Note To configure the enable password from the User Accounts pane (see the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1), change the password for the enable_15 user. The enable_15 user is always present in the User Accounts 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, and so on.

Step 4 In the Confirm Password field, reenter the password.

For security purposes, only asterisks appear in the password fields.

Step 5 To enable MS-CHAP authentication, check the **User authenticated using MSCHAP** check box.

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 MS-CHAPv1 or MS-CHAPv2.

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 by clicking 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 (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20), then this option lets the user use ASDM, SSH, Telnet, and the console port. If you also 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). See the "Configuring Command Authorization" section on page 52-23 for more information.
- **CLI login prompt for SSH, Telnet and console (no ASDM access)**—If you configure authentication for management access using the local database (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20), 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 (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20), 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 46-23.

Step 9 Click **Apply**.

The user is added to the local ASA database, and the changes are saved to the running configuration.

Tip You can search for specific text in each column of the Configuration > Device Management > Users/AAA > User Accounts pane. Enter the specific text that you want to locate in the Find box, then click the **Up** or **Down** arrow. You can also use the asterisk ("*") and question mark ("?") as wild card characters in the text search.

Configuring VPN Policy Attributes for a User

Prerequisites

This procedure describes how to edit an existing user. To add a user select **Configuration > Remote Access VPN > AAA/Local Users > Local Users** and click **Add**. For more information see "Adding a User Account to the Local Database" section on page 46-20.

Guidelines

By default, the **Inherit** check box is checked for each setting on the Edit User Account screen, which means that the user account inherits the value of that setting from the default group policy, DfltGrpPolicy.

To override each setting, uncheck the **Inherit** check box, and enter a new value. The detailed steps that follow describe each of the settings on the Edit User Account screen.

Detailed Steps

Step 1 Start ASDM and select **Configuration > Remote Access VPN > AAA/Local Users > Local Users**.

Step 2 Select the user you want configure and click **Edit**.

The Edit User Account screen opens.

Step 3 In the left-hand pane, click **VPN Policy**.

Step 4 Specify a group policy for the user. The user policy will inherit the attributes of this group policy. If there are other fields on this screen that are set to **Inherit** the configuration from the Default Group Policy, the attributes specified in this group policy will take precedence over those set in the Default Group Policy.

Step 5 Specify which tunneling protocols are available for the user, or whether the value is inherited from the group policy. Check the desired **Tunneling Protocols** check boxes to choose the VPN tunneling protocols that are available for use. Only the selected protocols are available for use. The choices are as follows:

- 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 shares (web-enabled), e-mail, and other TCP-based applications from almost any computer that can reach HTTPS Internet sites.
- The 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.

- IPsec IKEv1—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 Cisco VPN client-to-LAN connections can use IPsec IKEv1.
- IPsec IKEv2—Supported by the AnyConnect Secure Mobility Client. AnyConnect connections using IPsec with IKEv2 provide advanced features such as software updates, client profiles, GUI localization (translation) and customization, Cisco Secure Desktop, and SCEP proxy.
- 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 ASA and private corporate networks.

Note If no protocol is selected, an error message appears.

Step 6 Specify which filter (IPv4 or IPv6) 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 ASA, based on criteria such as source address, destination address, and protocol. To configure filters and rules, choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > General > More Options > Filter.**

Click **Manage** to display the ACL Manager pane, on which you can add, edit, and delete ACLs and ACEs.

Step 7 Specify whether to inherit the Connection Profile (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 users assigned group. If it is not, the ASA prevents the user from connecting. If the Inherit check box is not checked, the default value is None.

Step 8 Specify whether to inherit the Store Password on Client System setting from the group. Uncheck the **Inherit** check box to activate the Yes and No radio buttons. Click **Yes** to store the login password on the client system (potentially a less-secure option). Click **No** (the default) to require the user to enter the password with each connection. For maximum security, we recommend that you *not allow* password storage. This parameter has no effect on interactive hardware client authentication or individual user authentication for a VPN 3002.

Step 9 Specify an Access Hours policy to apply to this user, create a new access hours policy for the user, or leave the Inherit box checked. The default value is Inherit, or, if the Inherit check box is not checked, the default value is Unrestricted.

Click **Manage** to open the Add Time Range dialog box, in which you can specify a new set of access hours.

Step 10 Specify the number of simultaneous logins by the user. The Simultaneous Logins 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.

Step 11 Specify the **maximum connection time** for the 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, should we all be so lucky). To allow unlimited connection time, check the **Unlimited** check box (the default).

Step 12 Specify the Idle Timeout for the user in minutes. If there is no communication activity on the connection by this user 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 13 Configure the Session Alert Interval. If you uncheck the Inherit check box, the Default check box is checked automatically. This sets the session alert interval to 30 minutes. If you want to specify a new value, uncheck the Default check box and specify a session alert interval from 1 to 30 minutes in the minutes box.

Step 14 Configure the Idle Alert Interval. If you uncheck the Inherit check box, the Default check box is checked automatically. This sets the idle alert interval to 30 minutes. If you want to specify a new value, uncheck the Default check box and specify a session alert interval from 1 to 30 minutes in the minutes box.

Step 15 To set a dedicated IPv4 address for this user, enter an IPv4 address and subnet mask in the Dedicated IPv4 Address (Optional) area.

Step 16 To set a dedicated IPv6 address for this user, enter an IPv6 address with an IPv6 prefix in the Dedicated IPv6 Address (Optional) field. The IPv6 prefix indicates the subnet on which the IPv6 address resides.

Step 17 To configure clientless SSL settings, in the left-hand pane, click **Clientless SSL VPN**. To override each setting, uncheck the **Inherit** check box, and enter a new value.

Step 18 Click **Apply**.

The changes are saved to the running configuration.

Authenticating Users with a Public Key for SSH

Users can authenticate with a public key for SSH. The public key can be hashed or not hashed.

To authenticate with a public key for SSH, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Device Management > Users/AAA > User Accounts**.

Step 2 Select a user from the list, then click **Edit**.

The Edit User Account dialog box appears.

Step 3 Click **Public Key Authentication** in the navigation pane.

Step 4 If you want to hash the public key, check the **Key is hashed** check box. To not have the public key hashed, leave this check box unchecked.

If the public key is hashed, the value of the public key must have been previously hashed with SHA-256 and be 32 bytes long, with each byte separated by a colon (for parsing purposes).

If the public key is not hashed, the value of the key must be a Base 64 encoded public key that is generated by SSH key generation software that can generate SSH-RSA raw keys (that is, with no certificates). After you submit the Base 64 encoded public key, that key is then hashed via SHA-256 and the corresponding 32-byte hash is used for all further comparisons.

Step 5 Enter the public key.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration changes.

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 ASA when requiring user authentication from TACACS+ or RADIUS servers. This text is primarily for cosmetic purposes and appears above the username and password prompts that users see when they log in.

If you do not specify an authentication prompt, users 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, enter text in the Prompt field to add as a message to appear above the username and password prompts that users see when they log in.

The following table shows the allowed character limits 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 ASA displays the User accepted message text, if specified, to the user; otherwise, the ASA 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.

Testing Server Authentication and Authorization

To determine whether the ASA 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 in which the server resides.

The row is highlighted in the table.

Step 2 From the Servers in the Selected Group table, click the server that you want to test.

The row is highlighted in the table.

Step 3 Click **Test**.

The Test AAA Server dialog box appears for the selected server.

Step 4 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 5 In the Username field, enter a username.

Step 6 If you are testing authentication, in the Password field, enter the password for the username.

Step 7 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, ASDM displays an error message.

Monitoring AAA Servers

To monitor AAA servers, see the following panes:

Path	Purpose
Monitoring > Properties > AAA Servers	Shows the configured AAA server statistics.
Monitoring > Properties > AAA Servers	Shows the AAA server running configuration.
Choose **Tools > Command Line Interface**, enter the **show running-config all ldap attribute-map** command, then press **Send**.	Shows all LDAP attribute maps in the running configuration.
Choose **Tools > Command Line Interface**, enter the **show running-config zonelabs-integrity** command, then press **Send**.	Shows the Zone Labs Integrity server configuration.
Choose **Tools > Command Line Interface**, enter the **show ad-groups** *name* [**filter** *string*] command, then press **Send**.	Applies only to AD servers using LDAP, and shows groups that are listed on an AD server.

Additional References

For additional information related to implementing LDAP mapping, see the "RFCs" section on page 46-28.

RFCs

RFC	Title
2138	*Remote Authentication Dial In User Service (RADIUS)*
2139	*RADIUS Accounting*
2548	*Microsoft Vendor-specific RADIUS Attributes*
2868	*RADIUS Attributes for Tunnel Protocol Support*

Feature History for AAA Servers

Table 46-2 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 46-2 Feature History for AAA Servers

Feature Name	Platform Releases	Feature Information
AAA Servers	7.0(1)	AAA Servers describe support for AAA and how to configure AAA servers and the local database. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map Configuration > Device Management > Users/AAA > User Accounts Configuration > Device Management > Users/AAA > Authentication Prompt
Key vendor-specific attributes (VSAs) sent in RADIUS access request and accounting request packets from the ASA	8.4(3)	Four New VSAs—Tunnel Group Name (146) and Client Type (150) are sent in RADIUS access request packets from the ASA. Session Type (151) and Session Subtype (152) are sent in RADIUS accounting request packets from the ASA. All four attributes are sent for all accounting request packet types: Start, Interim-Update, and Stop. The RADIUS server (for example, ACS and ISE) can then enforce authorization and policy attributes or use them for accounting and billing purposes.

Configuring the Identity Firewall

This chapter describes how to configure the ASA for the Identity Firewall. The chapter includes the following sections:

- Information About the Identity Firewall, page 47-1
- Licensing for the Identity Firewall, page 47-7
- Guidelines and Limitations, page 47-7
- Prerequisites, page 47-8
- Configuring the Identity Firewall, page 47-9
- Monitoring the Identity Firewall, page 47-16
- Feature History for the Identity Firewall, page 47-18

Information About the Identity Firewall

This section includes the following topics:

- Overview of the Identity Firewall, page 47-1
- Architecture for Identity Firewall Deployments, page 47-2
- Features of the Identity Firewall, page 47-3
- Deployment Scenarios, page 47-4

Overview of the Identity Firewall

In an enterprise, users often need access to one or more server resources. Typically, a firewall is not aware of the users' identities and, therefore, cannot apply security policies based on identity. To configure per-user access policies, you must configure a user authentication proxy, which requires user interaction (a user name/password query).

The Identity Firewall in the ASA provides more granular access control based on users' identities. You can configure access rules and security policies based on user names and user groups name rather than through source IP addresses. The ASA applies the security policies based on an association of IP addresses to Windows Active Directory login information and reports events based on the mapped user names instead of network IP addresses.

The Identity Firewall integrates with Microsoft Active Directory in conjunction with an external Active Directory (AD) Agent that provides the actual identity mapping. The ASA uses Windows Active Directory as the source to retrieve the current user identity information for specific IP addresses and allows transparent authentication for Active Directory users.

Identity-based firewall services enhance the existing access control and security policy mechanisms by allowing users or groups to be specified in place of source IP addresses. Identity-based security policies can be interleaved without restriction between traditional IP address based rules.

The key benefits of the Identity Firewall include:

- Decoupling network topology from security policies
- Simplifying the creation of security policies
- Providing the ability to easily identify user activities on network resources
- Simplify user activity monitoring

Architecture for Identity Firewall Deployments

The Identity Firewall integrates with Window Active Directory in conjunction with an external Active Directory (AD) Agent that provides the actual identity mapping.

The identity firewall consists of three components:

- **ASA**
- **Microsoft Active Directory**

  Though Active Directory is part of the Identity Firewall on the ASA, they are managed by Active Directory administrators. The reliability and accuracy of the data depends on data in Active Directory.

  Supported versions include Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 servers.

- **Active Directory (AD) Agent**

  The AD Agent runs on a Windows server. Supported Windows servers include Windows 2003, Windows 2008, and Windows 2008 R2.

Note Windows 2003 R2 is not supported for the AD Agent server.

Figure 47-1 Identity Firewall Components

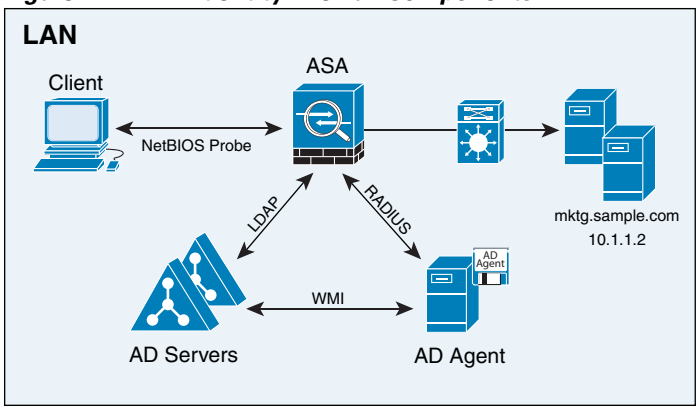

1	**On the ASA**: Configure local user groups and Identity Firewall policies.	**4**	**Client <-> ASA**: The client logs onto the network through Microsoft Active Directory. The AD Server authenticates users and generates user logon security logs. Alternatively, the client can log onto the network through a cut-through proxy or by using VPN.
2	**ASA <-> AD Server**: The ASA sends an LDAP query for the Active Directory groups configured on the AD Server. The ASA consolidates local and Active Directory groups and applies access rules and MPF security policies based on user identity.	**5**	**ASA <-> Client**: Based on the policies configured on the ASA, it grants or denies access to the client. If configured, the ASA probes the NetBIOS of the client to pass inactive and no-response users.
3	**ASA <-> AD Agent**: Depending on the Identity Firewall configuration, the ASA downloads the IP-user database or sends a RADIUS request to the AD Agent querying the user's IP address. The ASA forwards the new mappings learned from web authentication and VPN sessions to the AD Agent.	**6**	**AD Agent <-> AD Server**: Periodically or on-demand, the AD Agent monitors the AD Server security event log file via WMI for client login and logoff events. The AD Agent maintains a cache of user ID and IP address mappings. and notifies the ASA of changes. The AD Agent sends logs to a syslog server.

Features of the Identity Firewall

The Identity Firewall has the following key features.

Flexibility

- The ASA can retrieve user identity and IP address mappings from the AD Agent by querying the AD Agent for each new IP address or by maintaining a local copy of the entire user identity and IP address database.
- Supports host group, subnet, or IP address for the destination of a user identity policy.
- Supports a fully qualified domain name (FQDN) for the source and destination of a user identity policy.

- Supports the combination of 5-tuple policies with ID-based policies. The identity-based feature works in tandem with existing 5-tuple solution.
- Supports usage with IPS and Application Inspection policies.
- Retrieves user identity information from remote access VPN, AnyConnect VPN, L2TP VPN and cut-through proxy. All retrieved users are populated to all ASA devices connected to the AD Agent.

Scalability

- Each AD Agent supports 100 ASA devices. Multiple ASA devices are able to communicate with a single AD Agent to provide scalability in larger network deployments.
- Supports 30 Active Directory servers provided the IP address is unique among all domains.
- Each user identity in a domain can have up to 8 IP addresses.
- Supports up to 64,000 user identity-IP address mappings in active ASA policies for ASA 5500 Series models. This limit controls the maximum users who have policies applied. The total users are the aggregated users configured on all different contexts.
- Supports up to 1024 user identity-IP address mappings in active ASA policies for the ASA 5505.
- Supports up to 256 user groups in active ASA policies.
- A single rule can contain one or more user groups or users.
- Supports multiple domains.

Availability

- The ASA retrieves group information from Active Directory and falls back to web authentication for IP addresses that the AD Agent cannot map a source IP address to a user identity.
- The AD Agent continues to function when any of the Active Directory servers or the ASA are not responding.
- Supports configuring a primary AD Agent and a secondary AD Agent on the ASA. If the primary AD Agent stops responding, the ASA can switch to the secondary AD Agent.
- If the AD Agent is unavailable, the ASA can fall back to existing identity sources such as cut through proxy and VPN authentication.
- The AD Agent runs a watchdog process that automatically restarts its services when they are down.
- Allows a distributed IP address/user mapping database among ASA devices.

Deployment Scenarios

You can deploy the components of the Identity Firewall in the following ways depending on your environmental requirement.

As shown in Figure 47-2, you can deploy the components of the Identity Firewall to allow for redundancy. Scenario 1 shows a simple installation without component redundancy.

Scenario 2 also shows a simple installation without redundancy. However, in that deployment scenario, the Active Directory server and AD Agent are co-located on one Windows server.

Figure 47-2 Deployment Scenario without Redundancy

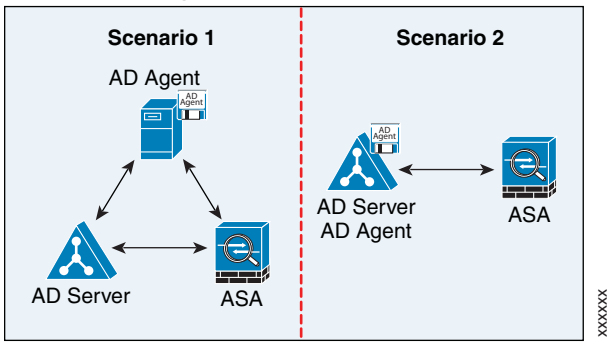

As shown in Figure 47-3, you can deploy the Identity Firewall components to support redundancy. Scenario 1 shows a deployment with multiple Active Directory servers and a single AD Agent installed on a separate Windows server. Scenario 2 shows a deployment with multiple Active Directory servers and multiple AD Agents installed on separate Windows servers.

Figure 47-3 Deployment Scenario with Redundant Components

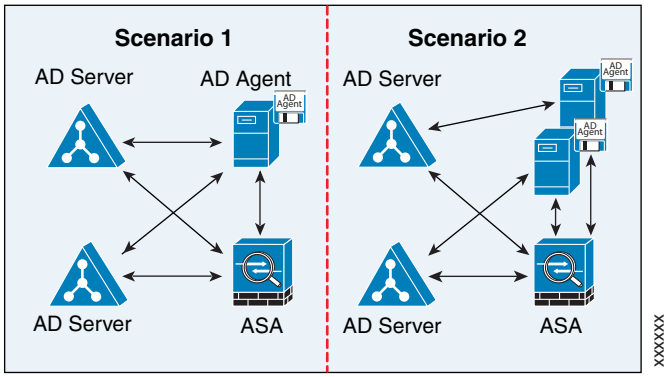

As shown in Figure 47-4, all Identity Firewall components—Active Directory server, the AD Agent, and the clients—are installed and communicate on the LAN.

Figure 47-4 LAN -based Deployment

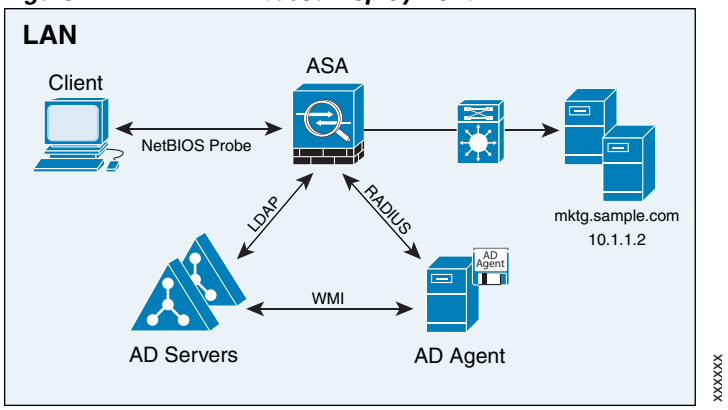

Figure 47-5 shows a WAN-based deployment to support a remote site. The Active Directory server and the AD Agent are installed on the main site LAN. The clients are located at a remote site and connect to the Identity Firewall components over a WAN.

Figure 47-5 WAN-based Deployment

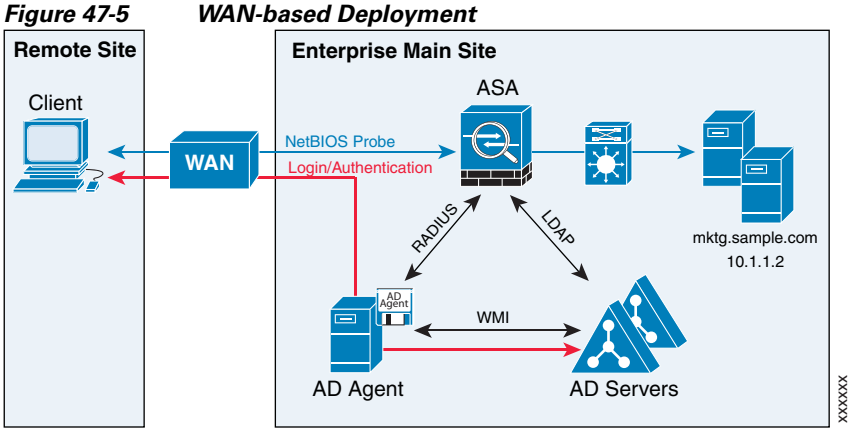

Figure 47-6 also shows a WAN-based deployment to support a remote site. The Active Directory server is installed on the main site LAN. However, the AD Agent is installed and access by the clients at the remote site. The remote clients connect to the Active Directory servers at the main site over a WAN.

Figure 47-6 WAN-based Deployment with Remote AD Agent

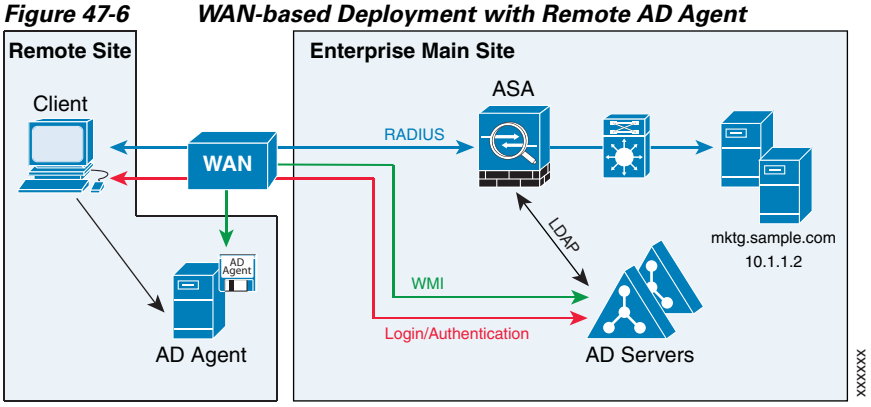

Figure 47-7 shows an expanded remote site installation. An AD Agent and Active Directory servers are installed at the remote site. The clients access these components locally when logging into network resources located at the main site. The remote Active Directory server must synchronize its data with the central Active Directory servers located at the main site.

Figure 47-7 WAN-based Deployment with Remote AD Agent and AD Servers

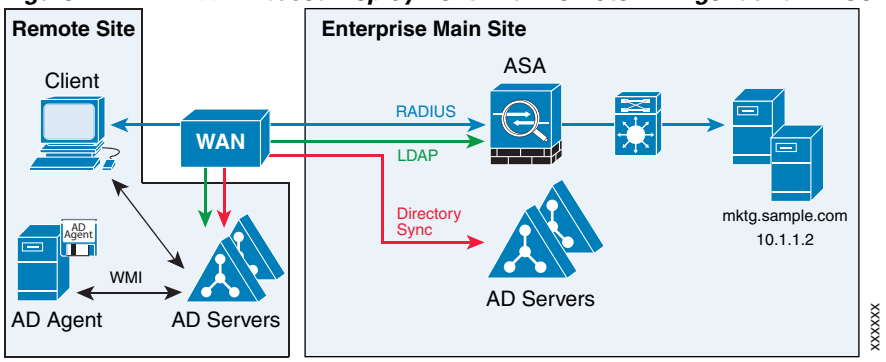

Licensing for the Identity Firewall

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 modes.

Failover Guidelines

The Identity Firewall supports user identity-IP address mappings and AD Agent status replication from active to standby when stateful failover is enabled. However, only user identity-IP address mappings, AD Agent status, and domain status are replicated. User and user group records are not replicated to the standby ASA.

When failover is configured, the standby ASA must also be configured to connect to the AD Agent directly to retrieve user groups. The standby ASA does not send NetBIOS packets to clients even when the NetBIOS probing options are configured for the Identity Firewall.

When a client is determined as inactive by the active ASA, the information is propagated to the standby ASA. User statistics are not propagated to the standby ASA.

When you have failover configured, you must configure the AD Agent to communicate with both the active and standby ASA devices. See the *Installation and Setup Guide for the Active Directory Agent* for the steps to configure the ASA on the AD Agent server.

IPv6 Guidelines

- Supports IPv6.

  The AD Agent supports endpoints with IPv6 addresses. It can receive IPv6 addresses in log events, maintain them in its cache, and send them through RADIUS messages.

- NetBIOS over IPv6 is not supported

Additional Guidelines and Limitations

- A full URL as a destination address is not supported.
- For NetBIOS probing to function, the network between the ASA, AD Agent, and clients must support UDP-encapsulated NetBIOS traffic.
- MAC address checking by the Identity Firewall does not work when intervening routers are present. Users logged onto clients that are behind the same router have the same MAC addresses. With this implementation, all the packets from the same router are able to pass the check, because the ASA is unable to ascertain to the actual MAC addresses behind the router.
- The following ASA features do not support using the identity-based object and FQDN in an extended ACL:
  - route-map
  - Crypto map
  - WCCP
  - NAT
  - group-policy (except VPN filter)
  - DAP

Prerequisites

Before configuring the Identity Firewall in the ASA, you must meet the prerequisites for the AD Agent and Microsoft Active Directory.

AD Agent

The AD Agent must be installed on a Windows server that is accessible to the ASA. Additionally, you must configure the AD Agent to obtain information from the Active Directory servers. Configure the AD Agent to communicate with the ASA.

Supported Windows servers include Windows 2003, Windows 2008, and Windows 2008 R2.

Note Windows 2003 R2 is not supported for the AD Agent server.

For the steps to install and configure the AD Agent, see the *Installation and Setup Guide for the Active Directory Agent*.

Before configuring the AD Agent in the ASA, obtain the secret key value that the AD Agent and the ASA use to communicate. This value must match on both the AD Agent and the ASA.

Microsoft Active Directory

Microsoft Active Directory must be installed on a Windows server and accessible by the ASA. Supported versions include Windows 2003, 2008, and 2008 R2 servers.

Before configuring the Active Directory server on the ASA, create a user account in Active Directory for the ASA.

Additionally, the ASA sends encrypted log in information to the Active Directory server by using SSL enabled over LDAP. SSL must be enabled on the Active Directory server. See the documentation for Microsft Active Diretory for the steps to enable SSL for Active Directory.

Note Before running the AD Agent Installer, you must install the following patches on every Microsoft Active Directory server that the AD Agent monitors. These patches are required even when the AD Agent is installed directly on the domain controller server. See the *README First for the Cisco Active Directory Agent*.

Configuring the Identity Firewall

This section contains the following topic:

- Task Flow for Configuring the Identity Firewall, page 47-9
- Configuring the Active Directory Domain, page 47-10
- Configuring Active Directory Server Groups, page 47-10
- Configuring Active Directory Server Groups, page 47-10
- Configuring Active Directory Agent Groups, page 47-12
- Configuring Identity Options, page 47-12
- Configuring Identity-Based Security Policy, page 47-15

Task Flow for Configuring the Identity Firewall

To configure the Identity Firewall, perform the following tasks:

Step 1 Configure the Active Directory domain in the ASA.

See Configuring the Active Directory Domain, page 47-10 and Configuring Active Directory Server Groups, page 47-10.

See also Deployment Scenarios, page 47-4 for the ways in which you can deploy the Active Directory servers to meet your environment requirements.

Step 2 Configure the AD Agent in ASA.

See Configuring Active Directory Server Groups, page 47-10 and Configuring Active Directory Agent Groups, page 47-12.

See also Deployment Scenarios, page 47-4 for the ways in which you can deploy the AD Agents to meet your environment requirements.

Step 3 Configure Identity Options.

See Configuring Identity Options, page 47-12.

Step 4 Configure Identity-based Security Policy.

After AD domain and AD-Agent are configured, you can create identity-based object groups and ACLs for use in many features. See Configuring Identity-Based Security Policy, page 47-15.

Configuring the Active Directory Domain

Active Directory domain configuration on the ASA is required for the ASA to download Active Directory groups and accept user identities from specific domains when receiving IP-user mapping from the AD Agent.

Prerequisites

- Active Directory server IP address
- Distinguished Name for LDAP base dn
- Distinguished Name and password for the Active Directory user that the Identity Firewall uses to connect to the Active Directory domain controller

To configure the Active Directory domain, perform the following steps:

Step 1 Choose **Configuration > Firewall > Identity Options**. The Identity Options pane appears.

Step 2 If necessary, check the Enable User Identity check box to enable the feature.

Step 3 In the Domains section, click **Add** or select a domain from the list and click **Edit**. The Domain dialog box appears.

Step 4 In the Domain NETBIOS Name field, enter a name up to 32 characters consisting of [a-z], [A-Z], [0-9], [!@#$%^&()-_=+[]{};,.] except '.' and ' ' at the first character. If the domain name contains a space, you must enclose that space character in quotation marks. The domain name is not case sensitive.

When you edit the name of an existing domain, the domain name associated with existing users and user groups is not changed.

Step 5 From the AD Server Group list, select the Active Directory servers to associate with this domain or click **Manage** to add a new server group to the list. See Configuring Active Directory Server Groups, page 47-10.

Step 6 Click **OK** to save the domain settings.

What to Do Next

Configure Active Directory server groups. See Configuring Active Directory Server Groups, page 47-10.

Configure AD Agents. See Configuring Active Directory Server Groups, page 47-10 and Configuring Active Directory Agent Groups, page 47-12.

Configuring Active Directory Server Groups

To configure the Active Directory server group, perform the following steps:

Step 1 Choose **Configuration > Firewall > Identity Options > Add > Manage**. The Configure Active Directory Server Groups dialog appears.

Step 2 To add an Active Directory server group for the Identify Firewall, click **Add** for the Active Directory Server Groups table. The Add Active Directory Server Group dialog box appears. See the "Configuring AAA Server Groups" section on page 46-11.

Step 3 To add servers to an Active Directory server group, select the group from the Active Directory Server Groups list and click **Add** for the Servers in the Selected Group table. The Add Active Directory Server dialog box appears. See the "Adding a Server to a Group" section on page 46-12.

Step 4 Click **OK** to save the settings.

What to Do Next

Configure AD Agents. See Configuring Active Directory Agents, page 47-11 and Configuring Active Directory Agent Groups, page 47-12.

Configuring Active Directory Agents

Periodically or on-demand, the AD Agent monitors the Active Directory server security event log file via WMI for user login and logoff events. The AD Agent maintains a cache of user ID and IP address mappings. and notifies the ASA of changes.

Requirement

- AD agent IP address
- Shared secret between ASA and AD agent

To configure the AD Agents, perform the following steps:

Step 1 Open the **Configuration > Firewall > Identity Options** pane.

Step 2 If necessary, check the Enable User Identity check box to enable the feature.

Step 3 In the Active Directory Agent section, click **Manage**.

The Configure Active Directory Agents dialog box appears.

Step 4 To add an AD Agent, click the **Add** button.

OR

Select an agent group from the list and click **Edit**.

See Configuring Active Directory Agent Groups, page 47-12.

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

What to Do Next

Configure AD Agent groups. See Configuring Active Directory Agent Groups, page 47-12.

Configure access rules for the Identity Firewall. See Configuring Identity-Based Security Policy, page 47-15.

Configuring Active Directory Agent Groups

Configure the primary and secondary AD Agents for the AD Agent Server Group. When the ASA detects that the primary AD Agent is not responding and a secondary agent is specified, the ASA switches to secondary AD Agent. The Active Directory server for the AD agent uses RADIUS as the communication protocol; therefore, you should specify a key attribute for the shared secret between ASA and AD Agent.

To configure the AD Agent Groups, perform the following steps:

Step 1 From the Configure Active Directory Agents dialog, click **Add**. The Add Active Directory Agent Group dialog box appears.

Step 2 Enter a name for the AD Agent group.

Step 3 From the Primary Active Directory Agent section, specify the interface on which the ASA listens for traffic from the AD Agent server, and enter the FQDN of the server or IP address.

Step 4 In the Primary Active Directory Agent section, enter a timeout interval and the retry interval for the attempts that the ASA will continue to contact the AD Agent when it is not responding.

Step 5 Enter the shared secret key that is used between primary AD Agent and the ASA.

Step 6 From the Secondary Active Directory Agent section, specify the interface on which the ASA listens for traffic from the AD Agent server, and enter the FQDN of the server or IP address.

Step 7 In the Secondary Active Directory Agent section, enter a timeout interval and the retry interval for the the attempts that the ASA will continue to contact the AD Agent when it is not responding.

Step 8 Enter the shared secret key that is used between secondary AD Agent and the ASA.

Step 9 Click **OK** to save your changes.

What to Do Next

Configure access rules for the Identity Firewall. See Configuring Identity-Based Security Policy, page 47-15.

Configuring Identity Options

Use this pane to add or edit the Identity Firewall feature; select the Enable check box to enable the feature. By default, the Identity Firewall feature is disabled.

Prerequisites

Before configuring the identify options for the Identity Firewall, you must you must meet the prerequisites for the AD Agent and Microsoft Active Directory. See Prerequisites, page 47-8 the requirements for the AD Agent and Microsoft Active Directory installation.

To configure the Identity Options for the Identity Firewall, perform the following steps:

Step 1 Choose **Configuration > Firewall > Identity Options**. The Identity Option pane appears.

Step 2 If necessary, check the Enable User Identity check box to enable the feature.

Step 3 To add a domain for the Identity Firewall, click **Add** by the Domains table. The Add Domain dialog box appears. See Configuring the Active Directory Domain, page 47-10.

Step 4 For domains already been added to the Domains list, check whether to disable rules when the domain is down because the Active Directory domain controller is not responding.

When a domain is down and this option is checked for that domain, the ASA disables the user identity rules associated with the users in that domain. Additionally, the status of all user IP addresses in that domain are marked as disabled in the **Monitoring > Properties > Identity > Users** pane.

Step 5 From the Default Domain drop-down list, select the default domain for the Identity Firewall.

The default domain is used for all users and user groups when a domain has not been explicitly configured for those users or groups. When a default domain is not specified, the default domain for users and groups is LOCAL.

Additionally, the Identity Firewall uses the LOCAL domain for all locally defined user groups or locally defined users (users who log in and authenticate by using a VPN or web portal).

Note The default domain name you select must match the NetBIOS domain name configured on the Active Directory domain controller. If the domain name does not match, the AD Agent will incorrectly associate the user-IP mappings with the domain name you enter when configuring the ASA.
To view the NetBIOS domain name, open the Active Directory user event security log in any text editor.

For multiple context modes, you can set a default domain name for each context, as well as within the system execution space.

Step 6 In the Active Directory Agent section, select the AD Agent group from the drop-down list. To add AD Agent groups, click **Manage**. See Configuring Active Directory Agents, page 47-11.

Step 7 In the Hello Timer field, enter a number between 10 to 65535 seconds.

The hello timer between the ASA and the AD Agent defines how frequently the ASA exchanges hello packets. The ASA uses the hello packet to obtain ASA replication status (in-sync or out-of-sync) and domain status (up or down). If the ASA does not receive a response from the AD Agent, it resends a hello packet after the specified interval.

Specify the number of times that the ASA will continue to send hello packets to the AD Agent. By default, the number of seconds is set to 30 and the retry times is set to 5.

Step 8 In the Poll Group Timer field, enter the number of hours that the ASA uses to query the DNS server to resolve fully-qualified domain names (FQDN). By default, the poll timer is set to 4 hours.

Step 9 In the Retrieve User Information, select an option from the list:

- On Demand—specifies that the ASA retrieve the user mapping information of an IP address from the AD Agent when the ASA receives a packet that requires a new connection and the user of its source IP address is not in the user-identity database.
- Full Download—specifies that the ASA send a request to the AD Agent to download the entire IP-user mapping table when the ASA starts and then to receive incremental IP-user mapping when users log in and log out.

Note Selecting On Demand has the benefit of using less memory as only users of received packets are queried and stored.

Step 10 In the Error Conditions section, select whether to disable rules in the AD Agent is not responding.

When the AD Agent is down and this option is selected, the ASA disables the user identity rules associated with the users in that domain. Additionally, the status of all user IP addresses in that domain are marked as disabled in **Monitoring > Properties > Identity > Users** pane.

Step 11 In the Error Conditions section, select whether to remove a user's IP address when the NetBIOS probe fails.

Selecting this option specifies the action when NetBIOS probing to a user is blocked (for example, the user client does not respond to a NetBIOS probe). The network connection might be blocked to that client or the client is not active. When this option is selected, the ASA disables the identity rules associated with that user's IP address.

Step 12 In the Error Conditions section, select whether to remove a user's MAC address when it is inconsistent with the IP address that the ASA has currently mapped to that MAC address. When this option is selected, the ASA disables the user identity rules associated with the specific user.

Step 13 In the Error Conditions section, select whether to track users that are not found.

Step 14 In the Users section, select the Idle Timeout option and enter a time in minutes from 1 minute to 65535. By default, the idle timeout is set to 60 minutes.

Enabling this option configures a timer when an active user is considered idle, meaning the ASA does not receive traffic from the user's IP address for more than the specified time. Once the timer expires, the user's IP address is marked inactive and removed from the local cached IP-user database and the ASA no longer notifies the AD Agent about that IP address. Existing traffic is still allowed to pass. When the Idle Timeout option is enabled, the ASA runs an inactive timer even when the NetBIOS Logout Probe is configured.

Note The Idle Timeout option does not apply to VPN or cut through proxy users.

Step 15 In the NetBIOS Logout Probe section, enable NetBIOS probing and set the probe timer (from1 to 65535 minutes) before a user's IP addresses is probed and the retry interval (from 1 to 256 retries) between retry probes.

Enabling this option configures how often the ASA probes the user host to determine whether the user client is still active. To minimize the NetBIOS packets, ASA only sends a NetBIOS probe to the client when the user has been idle for more than the specified number of minutes in the Idle Timeout minutes field.

Step 16 In the NetBIOS Logout Probe section, select an option from the User Name list:

- Match Any—As long as the NetBIOS response from the host contains the user name of the user assigned to the IP address, the user identity is be considered valid. Specifying this option requires that the host enabled the Messenger service and configured a WINS server.
- Exact Match—The user name of the user assigned to the IP address must be the only one in the NetBIOS response. Otherwise, the user identity of that IP address is considered invalid. Specifying this option requires that the host enabled the Messenger service and configured a WINS server.
- User Not Needed—As long as the ASA received a NetBIOS response from the host the user identity is considered valid.

Step 17 Click **Apply** to save the Identity Firewall Configuration.

What to Do Next

Configure the Active Directory domain and server groups. See Configuring the Active Directory Domain, page 47-10 and Configuring Active Directory Server Groups, page 47-10.

Configure AD Agents. See Configuring Active Directory Server Groups, page 47-10.

Configuring Identity-Based Security Policy

You can incorporate identity-based policy in many ASA features. Any feature that uses extended ACLs (other than those listed as unsupported in the "Guidelines and Limitations" section on page 47-7) can take advantage of identity firewall. You can now add user identity arguments to extended ACLs, as well as network-based parameters.

- To configure an extended ACL, see Chapter 19, "Adding an Extended Access Control List."
- To configure local user groups, which can be used in the ACL, see the "Configuring Local User Groups" section on page 25-7.

Features that can use identity include the following:

- Access rules—An access rule permits or denies traffic on an interface using network information. With identity firewall, you can now control access based on user identity. See Chapter 51, "Configuring Access Rules."
- AAA rules—An authentication rule (also known as "cut-through proxy") controls network access based on user. Because this function is very similar to an access rule + identity firewall, AAA rules can now be used as a backup method of authentication if a user's AD login expires. For example, for any user without a valid login, you can trigger a AAA rule. To ensure that the AAA rule is only triggered for users that do not have valid logins, you can specify special usernames in the extended ACL used for the access rule and for the AAA rule: None (users without a valid login) and Any (users with a valid login). In the access rule, configure your policy as usual for users and groups, but then include a rule that permits all None users; you must permit these users so they can later trigger a AAA rule. Then, configure a AAA rule that denies Any users (these users are not subject to the AAA rule, and were handled already by the access rule), but permits all None users:

```
access-list 100 ex permit ip user CISCO\xyz any any
access-list 100 ex deny ip user CISCO\abc any any
access-list 100 ex permit ip user NONE any any   ----> these users will match AAA rule
access-list 100 ex deny any any
access-group 100 in interface inside

access-list 200 ex deny ip user ANY any any -----> skips users who already logged in
access-list 200 ex permit user NONE any any
aaa authenticate match 200 inside user-identity
```

  See Chapter 53, "Configuring AAA Rules for Network Access."
- Cloud Web Security—You can control which users are sent to the Cloud Web Security proxy server. In addition, you can configure policy on the Cloud Web Security ScanCenter that is based on user groups that are included in ASA traffic headers sent to Cloud Web Security. See Chapter 75, "Configuring the ASA for Cisco Cloud Web Security."
- VPN filter—Although VPN does not support identity firewall ACLs in general, you can use configure the ASA to enforce identity-based access rules on VPN traffic. By default, VPN traffic is not subject to access rules. You can force VPN clients to abide by access rules that use an identity firewall ACL (**no sysopt connection permit-vpn** command). You can also use an identity firewall ACL with the VPN filter feature; VPN filter accomplishes a similar effect as allowing access rules in general.

- And many more...

Monitoring the Identity Firewall

This section contains the following topic:

- Monitoring AD Agents, page 47-16
- Monitoring Groups, page 47-16
- Monitoring Memory Usage for the Identity Firewall, page 47-17
- Monitoring Users for the Identity Firewall, page 47-17

Monitoring AD Agents

You can monitor the AD Agent component of the Identity Firewall.

Step 1 Open the **Monitoring > Properties > Identity > AD Agent**.

The Active Directory Agent Information pane appears.

Step 2 Click Refresh to update the data in the pane.

This pane displays the following information about the primary and secondary AD Agents:

- Status of the AD Agents
- Status of the domains
- Statistics for the AD Agents

Monitoring Groups

You can monitor the user groups configured for the Identity Firewall.

Step 1 Open the **Monitoring > Properties > Identity > Group**.

The Activated Groups pane appears.

Step 2 To display a list of the access rules using the selected group, click **Where used**.

Step 3 Click Refresh to update the data in the pane.

This pane displays displays the list of user groups in the following format:

domain\group_name

Monitoring Memory Usage for the Identity Firewall

You can monitor the memory usage that the Identity Firewall consumes on the ASA.

Step 1 Open the **Monitoring > Properties > Identity > Memory Usage**.

The Memory Usage of Identity Modules pane appears.

Step 2 Click Refresh to update the data in the pane.

This pane displays the memory usage in bytes of various modules in the Identity Firewall:

- Users
- Groups
- User Stats
- LDAP

  The ASA sends an LDAP query for the Active Directory groups configured on the Active Directory server. The Active Directory server authenticates users and generates user logon security logs.
- AD Agent
- Miscellaneous
- Total Memory Usage

Note How you configure the Identity Firewall to retrieve user information from the AD Agent impacts the amount of memory used by the feature. You specify whether the ASA uses **on demand** retrieval or **full download** retrieval. Selecting **On Demand** has the benefit of using less memory as only users of received packets are queried and stored. See Configuring Identity Options, page 47-12 for a description of these options.

Monitoring Users for the Identity Firewall

You can display information about all users contained in the IP-user mapping database used by the Identity Firewall.

Step 1 Open the **Monitoring > Properties > Identity > User**.

The Users in the User Database pane appears.

Note Active users are highlighted in green.

Step 2 To display additional information about an active user, select the user in the list and click **Details**. The Details button is enabled for active users only.

Step 3 To display a list of the access rules using the selected user, click **Where used**.

Step 4 Click Refresh to update the data in the pane.

This pane displays the following information for users:

domain\user_name	Status (active or inactive)	Connections	Minutes Idle

The default domain name can be the real domain name, a special reserved word, or LOCAL. The Identity Firewall uses the LOCAL domain name for all locally defined user groups or locally defined users (users who log in and authenticate by using a VPN or web portal). When default domain is not specified, the default domain is LOCAL.

The idle time is stored on a per user basis instead of per the IP address of a user.

If the option to disable rules when the Active Directory server is down and the domain is down, or the option to disable rules in the AD Agent is down and the AD Agent is down, all the logged on users have the status disabled. You configure these options in the Identity Options pane.

Alternatively, you can view statistics for users by accessing the Firewall Dashboard pane. The Firewall Dashboard pane lets you view important information about the traffic passing through your ASA. Choose **Home > Firewall Dashboard > Top 10 Users** tab in the Top Usage Status area.

The Top 10 Users tab displays data only when you have configured the Identity Firewall feature in the ASA, which includes configuring these additional components—Microsoft Active Directory and Cisco Active Directory (AD) Agent. See Configuring the Identity Firewall, page 47-9 for information.

Depending on which option you choose, the Top 10 Users tab shows statistics for received EPS packets, sent EPS packets, and sent attacks for the top 10 users. For each user (displayed as *domain\user_name*), the tab displays the average EPS packet, the current EPS packet, the trigger, and total events for that user.

Note The first three tabs in the Top Usage Status area display threat detection data and are unrelated to the Identity Firewall feature.

Feature History for the Identity Firewall

Table 47-1 lists the release history for this feature.

Table 47-1 Feature History for the Identity Firewall

Feature Name	Releases	Feature Information
Identity Firewall	8.4(2)	The Identity Firewall feature was introduced. We introduced or modified the following screens: Configuration > Firewall > Identity Options Configuration > Firewall > Objects > Local User Groups Monitoring > Properties > Identity

Configuring the ASA to Integrate with Cisco TrustSec

This chapter includes the following sections:

- Information About the ASA Integrated with Cisco TrustSec, page 48-1
- Licensing Requirements when Integrating the ASA with Cisco TrustSec, page 48-8
- Prerequisites for Integrating the ASA with Cisco TrustSec, page 48-8
- Guidelines and Limitations, page 48-9
- Configuring the ASA for Cisco TrustSec Integration, page 48-11
- Monitoring the ASA Integrated with Cisco TrustSec, page 48-17
- Feature History for the ASA-Cisco TrustSec Integration, page 48-21

Information About the ASA Integrated with Cisco TrustSec

This section includes the following topics:

- Information about Cisco TrustSec, page 48-1
- About SGT and SXP Support in Cisco TrustSec, page 48-2
- Roles in the Cisco TrustSec Solution, page 48-3
- Security Group Policy Enforcement, page 48-3
- How the ASA Enforces Security Group Based Policies, page 48-4
- About Speaker and Listener Roles on the ASA, page 48-6
- Features of the ASA-Cisco TrustSec Integration, page 48-6

Information about Cisco TrustSec

Traditionally, security features such as firewalls performed access control based on predefined IP addresses, subnets and protocols. However, with enterprises transitioning to borderless networks, both the technology used to connect people and organizations and the security requirements for protecting data and networks have evolved significantly. End points are becoming increasingly nomadic and users often utilize a variety of end points (for example, laptop versus desktop, smart phone, or tablet), which

means that a combination of user attributes plus end-point attributes provide the key characteristics, in addition to existing 5-tuple based rules, that enforcement devices, such as switches and routers with firewall features or dedicated firewalls, can reliably use for making access control decisions.

As a result, the availability and propagation of end point attributes or client identity attributes have become increasingly important requirements to enable security solutions across the customers' networks, at the access, distribution, and core layers of the network and in the data center to name but a few examples.

Cisco TrustSec provides an access-control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec solution, enforcement devices utilize a combination of user attributes and end-point attributes to make role-based and identity-based access control decisions. The availability and propagation of this information enables security solutions across networks at the access, distribution, and core layers of the network.

Implementing Cisco TrustSec into your environment has the following advantages:

- Provides a growing mobile and complex workforce with appropriate and more secure access from any device
- Lowers security risks by providing comprehensive visibility of who and what is connecting to the wired or wireless network
- Offers exceptional control over activity of network users accessing physical or cloud-based IT resources
- Reduces total cost of ownership through centralized, highly secure access policy management and scalable enforcement mechanisms

For information about Cisco TrustSec, see http://www.cisco.com/go/trustsec.

About SGT and SXP Support in Cisco TrustSec

In the Cisco TrustSec solution, security group access transforms a topology-aware network into a role-based network, thus enabling end-to-end policies enforced on the basis of role-based access-control (RBACL). Device and user credentials acquired during authentication are used to classify packets by security groups. Every packet entering the Cisco TrustSec cloud is tagged with an security group tag (SGT). The tagging helps trusted intermediaries identify the source identity of the packet and enforce security policies along the data path. An SGT can indicate a privilege level across the domain when the SGT is used to define a security group access list.

An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB), which happens with a RADIUS vendor-specific attribute. An SGT can be assigned statically to a particular IP address or to a switch interface. An SGT is passed along dynamically to a switch or access point after successful authentication.

The Security-group eXchange Protocol (SXP) is a protocol developed for Cisco TrustSec to propagate the IP-to-SGT mapping database across network devices that do not have SGT-capable hardware support to hardware that supports SGTs and security group access lists. SXP, a control plane protocol, passes IP-SGT mappings from authentication points (such as legacy access layer switches) to upstream devices in the network.

The SXP connections are point-to-point and use TCP as the underlying transport protocol. SXP uses the well known TCP port number 64999 when initiating a connection. Additionally, an SXP connection is uniquely identified by the source and destination IP addresses.

Roles in the Cisco TrustSec Solution

To provide identity and policy-based access enforcement, the Cisco TrustSec solution includes the functionality:

- **Access Requestor (AR)**: Access requestors are end-point devices that request access to protected resources in the network. They are primary subjects of the architecture and their access privilege depends on their Identity credentials.

  Access requestors include end-point devices such PCs, laptops, mobile phones, printers, cameras, and MACsec-capable IP phones.

- **Policy Decision Point (PDP)**: A policy decision point is responsible for making access control decisions. The PDP provides features such as 802.1x, MAB, and Web authentication. The PDP supports authorization and enforcement through VLAN, DACL, and security group access (SGACL/SXP/SGT).

  In the Cisco TrustSec solution, the Cisco Identity Services Engine (ISE) acts as the PDP. The Cisco ISE provides identity and access control policy functionality.

- **Policy Information Point (PIP)**: A policy information point is a source that provides external information (for example, reputation, location, and LDAP attributes) to policy decision points.

  Policy information points include devices such as Session Directory, Sensors IPS, and Communication Manager.

- **Policy Administration Point (PAP)**: A policy administration point defines and inserts policies into authorization system. The PAP acts as an identity repository, by providing Cisco TrustSec tag to user identity mapping and Cisco Trustsec tag to server resource mapping.

  In the Cisco TrustSec solution, the Cisco Secure Access Control System (a policy server with integrated 802.1x and SGT support) acts as the PAP.

- **Policy Enforcement Point (PEP)**: A policy enforcement point is the entity that carries out the decisions (policy rules and actions) made by the PDP for each AR. PEP devices learn identity information through the primary communication path that exists across networks. PEP devices learn the identity attributes of each AR from many sources, such as end-point agents, authorization servers, peer-enforcement devices, and network flows. In turn, PEP devices use SXP to propagate IP-SGT mappings to mutually-trusted peer devices across the network.

  Policy enforcement points include network devices such as Catalyst switches, routers, firewalls (specifically the ASA), servers, VPN devices, and SAN devices.

The ASA serves the role of the PEP in the identity architecture. Using SXP, the ASA learns identity information directly from authentication points and uses that to enforce identity-based policies.

Security Group Policy Enforcement

Security policy enforcement is based on security group name. An end-point device attempts to access a resource in the data center. Compared to traditional IP-based policies configured on firewalls, identity-based policies are configured based on user and device identities. For example, mktg-contractor is allowed to access mktg-servers; mktg-corp-users are allowed to access mktg-server and corp-servers.

The benefits of this type of deployment include:

- User group and Resource is defined and enforced using single object (SGT) – simplified policy management.

- User identity and resource identity are retained throughout the Cisco Trustsec capable switch infrastructure.

Figure 48-1 Security Group Name Based Policy Enforcement Deployment

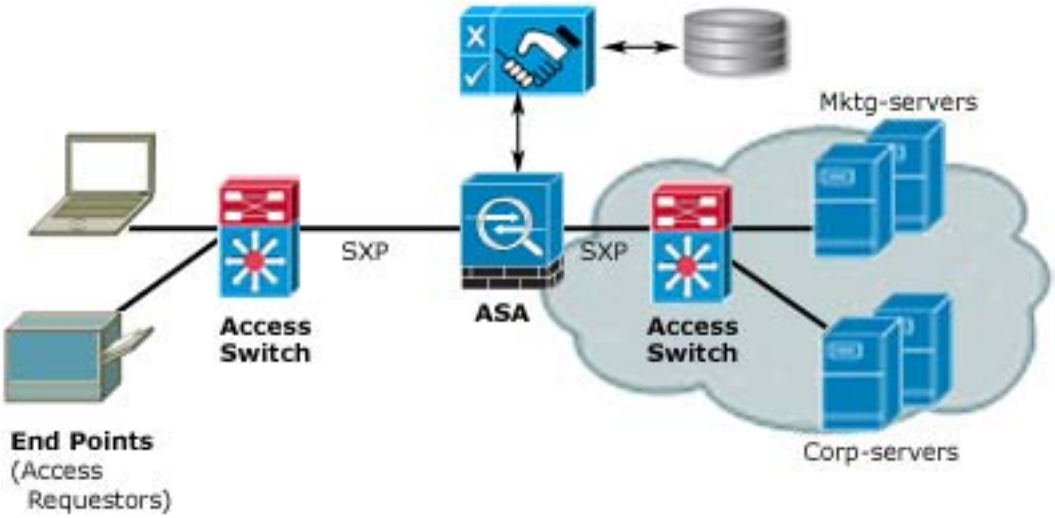

Implementing Cisco TrustSec allows for configuration of security policies supporting server segmentation.

- A pool of servers can be assigned an SGT for simplified policy management.
- The SGT information is retained within the infrastructure of Cisco Trustsec capable switches.
- The ASA can leverage the IP-SGT mapping for policy enforcement across the Cisco TrustSec domain.
- Deployment simplification is possible because 802.1x authorization for servers is mandatory.

How the ASA Enforces Security Group Based Policies

Note User-based security policies and security-group based policies, can coexist on the ASA. Any combination of network, user-based and security-group based attributes can be configured in an security policy. See Chapter 47, "Configuring the Identity Firewall" for information about configuring user-based security policies.

As part of configuring the ASA to function with Cisco TrustSec, you must import a Protected Access Credential (PAC) file from the ISE. Importing a Protected Access Credential (PAC) File, page 48-13.

Importing the PAC file to the ASA establishes a secure communication channel with the ISE. After the channel is established, the ASA initiates a PAC secure RADIUS transaction with the ISE and downloads Cisco TrustSec environment data; specifically, the ASA downloads the security group table. The security group table maps SGTs to security group names. Security group names are created on the ISE and provide user-friendly names for security groups.

The first time the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names contained in security policies configured on the ASA; then, the ASA activates those security policies locally. If the ASA is unable to resolve a security group name, it generates a system log message for the unknown security group name.

The following figure shows how a security policy is enforced in Cisco TrustSec.

Figure 48-2 Security Policy Enforcement

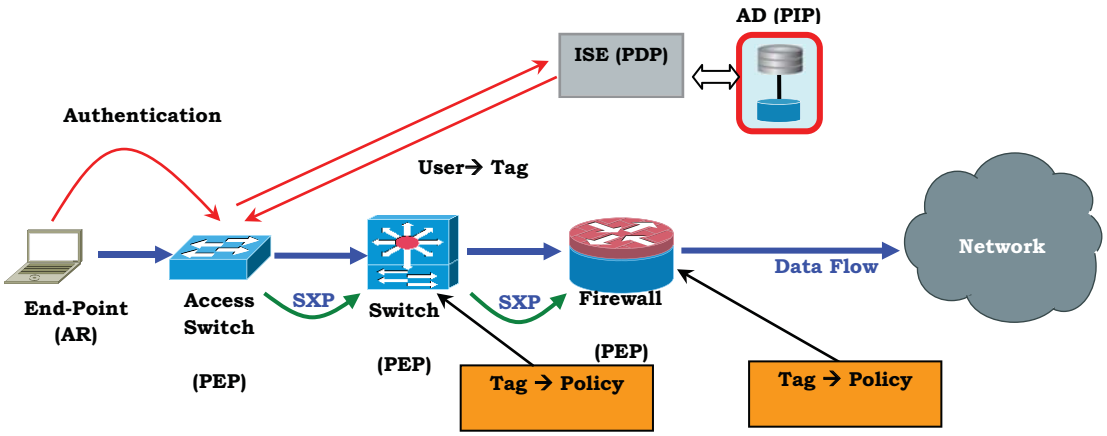

1. An end-point device connects to an access layer device directly or via remote access and authenticates with Cisco TrustSec.
2. The access layer device authenticates the end-point device with the ISE by using authentication methods such as 802.1X or web authentication. The end-point device passes role and group membership to classify the device into the appropriate security group.
3. The access layer device uses SXP to propagate the IP-SGT mapping to the upstream devices.
4. The ASA receives the packet. Using the IP-SGT mapping passed by SXP, the ASA looks up the SGTs for the source and destination IP addresses.

   If the mapping is new, the ASA records it in its local IP-SGT Manager database. The IP-SGT Manager database, which runs in the control plan, tracks IP-SGT mappings for each IPv4 or IPv6 address. The database records the source from which the mapping was learned. The peer IP address of the SXP connection is used as the source of the mapping. Multiple sources can exist for each IP-SGT mapping.

   If the ASA is configured as a Speaker, the ASA transmits all IP-SGT mappings to its SXP peers. See About Speaker and Listener Roles on the ASA, page 48-6.
5. If a security policy is configured on the ASA with that SGT or security group name, the ASA enforces the policy. (You can create security policies on the ASAthat contain SGTs or security group names. To enforce policies based on security group names, the ASA needs the security group table to map security group names to SGTs.)

   If the ASA cannot find a security group name in the security group table and it is included in a security policy, the ASA considers the security group name unknown and generates a system log message. When it becomes know after the ASA refreshes the security group table from the ISE, the ASA generates a system log message indicating that the security group name is known.

About Speaker and Listener Roles on the ASA

The ASA supports SXP to send and receive IP-SGT mappings to and from other network devices. Employing SXP allows security devices and firewalls to learn identity information from access switches without the need for hardware upgrades or changes. SXP can also be used to pass IP-SGT mappings from upstream devices (such as datacenter devices) back to the downstream devices. The ASA can receive information from both upstream and downstream directions.

When configuring an SXP connection on the ASA to an SXP peer, you must designate the ASA as a Speaker or a Listener for that connection so that it can exchange identity information:

- **Speaker mode**—configures the ASA so that it can forward all active IP-SGT mappings collected on the ASA to upstream devices for policy enforcement.
- **Listener mode**—configures the ASA so that it can receive IP-SGT mappings from downstream devices (SGT-capable switches) and use that information in creating policy definitions.

If one end of an SXP connection is configured as Speaker, then the other end must be configured as a Listener, and vice versa. If both devices on each end of an SXP connection are configured with the same role (either both as Speakers or both as Listeners), the SXP connection will fail and the ASA will generate a system log message.

Configuring the ASA to be both a Speaker and a Listener for an SXP connection can cause SXP looping, meanings that SXP data can be received by an SXP peer that originally transmitted it.

As part of configuring SXP on the ASA, you configure an SXP reconcile timer. After an SXP peer terminates its SXP connection, the ASA starts a hold down timer. Only SXP peers designated as Listener devices can terminate a connection. If an SXP peer connects while the hold down timer is running, the ASA starts the reconcile timer; then, the ASA updates the IP-SGT mapping database to learn the latest mappings.

Features of the ASA-Cisco TrustSec Integration

The ASA leverages Cisco TrustSec as part of its identity-based firewall feature. The integrating the ASA with Cisco TrustSec provides the following key features.

Flexibility

- The ASA can be configured as an SXP Speaker or Listener, or both.

  See About Speaker and Listener Roles on the ASA, page 48-6.
- The ASA supports SXP for IPv6 and IPv6 capable network devices.
- The ASA negotiates SXP versions with different SXP-capable network devices. SXP version negotiation eliminates the need for static configuration of versions.
- You can configure the ASA to refresh the security group table when the SXP reconcile timer expires and you can download the security group table on demand. When the security group table on the ASA is updated from the ISE, changes are reflected in the appropriate security policies.
- The ASA supports security policies based on security group names in the source or destination fields, or both. You can configure security policies on the ASA based on combinations of security groups, IP address, Active Directory group/user name, and FQDN.

Availability

- You can configure security group based policies on the ASA in Active/Active and Active/Standby configuration.

- The ASA can communicate with the ISE configured for high availability (HA).
- If the PAC file downloaded from the ISE expires on the ASA and the ASA cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until the ASA downloads an updated table.

Scalability

The ASA supports the following number of IP-SGT mapped entries:

Table 48-1 Capacity Numbers for IP-SGT Mappings

ASA Platform	Number of IP-SGT Mapped Entries
5505	250
5510	1000
5520	2500
5540	5000
5550	7500
5580-20	10,000
5580-40	20,000
5585-X with SSP-10	18,750
5585-X with SSP-20	25,000
5585-X with SSP-40	50,000
5585-X with SSP-60	100,000

The ASA supports the following number of SXP connections:

Table 48-2 SXP Connections

ASA Platform	Number of SXP TCP Connections
5505	10
5510	25
5520	50
5540	100
5550	150
5580-20	250
5580-40	500
5585-X with SSP-10	150
5585-X with SSP-20	250
5585-X with SSP-40	500
5585-X with SSP-60	1000

Licensing Requirements when Integrating the ASA with Cisco TrustSec

Model	License Requirement
All models	Base License.

Prerequisites for Integrating the ASA with Cisco TrustSec

Before configuring the ASA to integrate with Cisco TrustSec, you must perform the following prerequisites:

- Register the ASA with the ISE.
- Create a security group for the ASA on the ISE.
- Generate the PAC file on the ISE to import into the ASA.

Registering the ASA with the ISE

The ASA must be configured as a recognized Cisco TrustSec network device in the ISE before the ASA can successfully import a PAC file.

1. Log into the ISE.
2. Choose **Administration > Network Devices > Network Devices**.
3. Click **Add**.
4. Enter the IP address of the ASA.
5. When the ISE is being used for user authentication in the Cisco TrustSec solution, enter a shared secret in the Authentication Settings area.

   When you configure the AAA sever on the ASA, provide the shared secret you create here on the ISE. The AAA server on the ASA uses this shared secret to communicate with the ISE.
6. Specify a device name, device ID, password, and a download interval for the ASA. See the ISE documentation for the details to perform these tasks.

Creating a Security Group on the ISE

When configuring the ASA to communicate with the ISE, you specify a AAA server. When configuring the AAA server on the ASA, you must specify a server group.

The security group must be configured to use the RADIUS protocol.

1. Log into the ISE.
2. Choose **Policy > Policy Elements > Results > Security Group Access > Security Group**.
3. Add a security group for the ASA. (Security groups are global and not ASA specific.)

   The ISE creates an entry under Security Groups with a tag.
4. Under the Security Group Access section, configure a device ID credentials and password for the ASA.

Generating the PAC File

For information about the PAC file, see Importing a Protected Access Credential (PAC) File, page 48-13.

Before generating the PAC file, you must have registered the ASA with the ISE.

1. Log into the ISE.
2. Choose **Administration > Network Resources > Network Devices**.
3. From the list of devices, select the ASA device.
4. Under the Security Group Access (SGA), click **Generate PAC**.
5. To encrypt the PAC file, enter a password.

   The password (or encryption key) you enter to encrypt the PAC file is independent of the password that was configured on the ISE as part of the device credentials.

The ISE generates the PAC file. The ASA can import the PAC from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC does not have to reside on the ASA flash before you can import it.)

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

Clustering Guideline

Supported only on the master device in a clustering setting.

High Availability Guideline

Supports a list of servers via configuration. If the first server is unreachable, the ASA will try to contact the second server in the list, and so on. However, the server list downloaded as part of the Cisco TrustSec environment data is ignored.

Limitations

- The ASA can only be configured to interoperate in a single Cisco TrustSec domain.
- The ASA does not support static configuration of SGT-name mappings on the device.
- NAT is not supported in SXP messages.
- SXP conveys IP-SGT mappings to enforcement points in the network. If an access layer switch belongs to a different NAT domain than the enforcing point, the IP-SGT map it uploads is invalid and an IP-SGT mappings database lookup on the enforcement device will not yield valid results; therefore, the ASA cannot apply security group aware security policy on the enforcement device.

- When configuring the ASA to integrate with Cisco TrustSec, you configure a default password for the ASA to use for SXP connections, or you can elect to not use a password; however, connection specific passwords are not supported for SXP peers. The configured default SXP password should be consistent cross the deployment network.
- SXP connection loops can form when a device has bidirectional connections to a peer, or is part of a unidirectionally connected chain of devices. (The ASA can learn IP-DGT mappings for resources from the access layer in the data center. The ASA might need to propagate these tags to downstream devices.) SXP connection loops can cause unexpected behavior of SXP message transport. In cases where the ASA is configured to be a Speaker and Listener, an SXP connection loop can occur causing SXP data to be received by the peer that originally transmitted it.
- When changing the ASA local IP address, you must ensure that all SXP peers have updated their peer list. Likewise, if an SXP peers changes its IP addresses, you must ensure those changes are reflected on the ASA.
- Automatic PAC file provisioning is not supported. The ASA administrator must request the PAC file from the ISE administrative interface and import it to the ASA. For information about the PAC file, see Generating the PAC File, page 48-9 and Importing a Protected Access Credential (PAC) File, page 48-13.
- PAC files have expiration dates. You must import the updated PAC file before the current PAC file expires; otherwise, the ASA will not be able to retrieve environment data updates.
- When a security group changes on the ISE (for example, it is renamed or deleted), the ASA does not change the status of any ASA security policies that contain an SGT or security group name associated with the changed security group; however, the ASA generates a system log message to indicate that those security policies changed.

  See Refreshing Environment Data, page 48-16 for information about manually updating the security group table on the ASA to pick up changes from the ISE.
- The multicast types are not supported in ISE 1.0.
- An SXP connection stays in the initializing state among two SXP peers interconnected by the ASA; as shown in the following example:

  (SXP peer A) - - - - (ASA) - - - (SXP peer B)

  Therefore, when configuring the ASA to integrate with Cisco TrustSec, you must enable the no-NAT, no-SEQ-RAND, and MD5-AUTHENTICATION TCP options on the ASA. Create a TCP-state-bypass policy for traffic destined to SXP port TCP 64999 among the SXP peers. Apply the policy on the appropriate interfaces.

  For example, configure the ASA as shown in this sample configuration for a TCP-state-bypass policy:

```
access-list SXP-MD5-ACL extended permit tcp host <peerA> host <peerB> eq 64999
access-list SXP-MD5-ACL extended permit tcp host <peerB> host <peerA> eq 64999

tcp-map SXP-MD5-OPTION-ALLOW
 tcp-options range 19 19 allow

class-map SXP-MD5-CLASSMAP
 match access-list SXP-MD5-ACL

policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
policy-map global_policy
class SXP-MD5-CLASSMAP
  set connection random-sequence-number disable
  set connection advanced-options SXP-MD5-OPTION-ALLOW
```

```
  set connection advanced-options tcp-state-bypass
service-policy global_policy global
```

Configuring the ASA for Cisco TrustSec Integration

This section contains the following topics:

- Task Flow for Configuring the ASA to Integrate with Cisco TrustSec, page 48-11
- Configuring the AAA Server for Cisco TrustSec Integration, page 48-12
- Importing a Protected Access Credential (PAC) File, page 48-13
- Configuring the Security Exchange Protocol (SXP), page 48-14
- Adding an SXP Connection Peer, page 48-15
- Refreshing Environment Data, page 48-16
- Configuring the Security Policy, page 48-17

Task Flow for Configuring the ASA to Integrate with Cisco TrustSec

Prerequisite

Before configuring the ASA to integrate with Cisco TrustSec, you must meet the following prerequisites:

- Register the ASA with the ISE.
- Generate the PAC file on the ISE to import into the ASA.

See the "Prerequisites for Integrating the ASA with Cisco TrustSec" section on page 48-8 for information.

Task Flow in the ASA

To configure the ASA to integrate with Cisco TrustSec, perform the following tasks:

Step 1 Configure the AAA server.

See Configuring the AAA Server for Cisco TrustSec Integration, page 48-12.

Step 2 Import the PAC file from the ISE.

See Importing a Protected Access Credential (PAC) File, page 48-13.

Step 3 Enable and set the default values for SXP.

See Configuring the Security Exchange Protocol (SXP), page 48-14.

Step 4 Add SXP connection peers for the Cisco TrustSec architecture.

See Adding an SXP Connection Peer, page 48-15.

Step 5 As necessary, refresh environment data for the ASA integrated with Cisco TrustSec.

See Refreshing Environment Data, page 48-16.

Step 6 Configure the Security Policy.

See Configuring the Security Policy, page 48-17.

Configuring the AAA Server for Cisco TrustSec Integration

As part of configuring the ASA to integrate with Cisco TrustSec, you must configure the ASA so that it can communicate with the ISE.

See also the "Configuring AAA Server Groups" section on page 46-11 and the "Adding a Server to a Group" section on page 46-12 for more information.

Prerequisites

- The referenced server group must be configured to use the RADIUS protocol. If you add a non-RADIUS server group to the ASA, the feature configuration will fail.
- If the ISE is also used for user authentication, obtain the shared secret that was entered on the ISE when you registered the ASA with the ISE. Contact your ISE administrator if you do not have this information.

To configure the ASA to communicate with the ISE for Cisco TrustSec integration, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**. The Identity By TrustSec pane appears.

Step 2 To add a server group to the ASA, click **Manage** in the Server Group Setup area, The Configure AAA Server Group dialog box appears.

Step 3 In the AAA Server Group field, enter the name of the security group created on the ISE for the ASA.

The server group name you specify here must match the name of the security group created on the ISE for the ASA. If these two group names do not match, the ASA will not be able to communicate with the ISE. Contact your ISE administrator if you do not have this information.

Step 4 In the Protocol drop-down list, select RADIUS.

For information about completing the remaining fields in the AAA Server Group dialog box, see Configuring AAA Server Groups, page 46-11

Step 5 Click **OK**. The ASA adds the group to the list of AAA Server Groups.

Step 6 To add a server to a group, select the AAA sever group you just created and click **Add** in the Servers in the Selected Group area (lower pane). The Add AAA Server dialog box appears.

Step 7 In the Interface Name field, select the network interface where the ISE server resides.

Step 8 In the Server Name or IP Address field, enter the IP address of the ISE server.

For information about completing the remaining fields in the AAA Server dialog box, see See Adding a Server to a Group, page 46-12.

Step 9 Click **OK**. The ASA adds the ISE server to the list of AAA servers.

Step 10 Click **OK** to save the addition of the ISE server and server group for the integration with Cisco TrustSec.

The changes are saved to the running configuration.

Importing a Protected Access Credential (PAC) File

Importing the PAC file to the ASA establishes the connection with the ISE. After the channel is established, the ASA initiates a secure RADIUS transaction with the ISE and downloads Cisco TrustSec environment data; specifically, the ASA downloads the security group table. The security group table maps SGTs to security group names. Security group names are created on the ISE and provide user-friendly names for security groups.

More specifically, no channel is established prior to the radius transaction. The ASA simply initiates a radius transaction with the ISE using the PAC for authentication

Tip The PAC file contains a shared key that allows the ASA and ISE to secure the RADIUS transactions that occur between them. Given the sensitive nature of this key, it must be stored securely on the ASA.

When you import the PAC file, the file is converted to ASCII HEX format and sent to the ASA in non-interactive mode. After successfully importing the file, the ASA download Cisco TrustSec environment data from the ISE without requiring the device password configured in the ISE.

Prerequisites

- The ASA must be configured as a recognized Cisco TrustSec network device in the ISE before the ASA can generate a PAC file. The ASA can import any PAC file but it will only work on the ASA when the file was generated by a properly configured ISE. See Registering the ASA with the ISE, page 48-8.
- Obtain the password used to encrypt the PAC file when generating it on the ISE.

  The ASA requires this password to import and decrypt the PAC file.
- Access to the PAC file generated by the ISE. The ASA can import the PAC from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC does not have to reside on the ASA flash before you can import it.)
- The server group has been configured for the ASA.

Restrictions

- When the ASA is part of an HA configuration, you must import the PAC file to the primary ASA device.
- When the ASA is part of a clustering configuration, you must import the PAC to the master device.

To import a PAC file, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**. The Identity By TrustSec pane appears.

Step 2 Check the Enable Security Exchange Protocol checkbox to enable SXP.

Step 3 In the Server Group Setup area, click Import PAC. The Import PAC dialog box appears.

Step 4 In the Filename field, enter the path and filename for the PAC file by using one of the following formats:

- disk0: Path and filename on disk0
- disk1: Path and filename on disk1
- flash: Path and filename on flash

Step 5 In the Password field, enter the password used to encrypt the PAC file. The password is independent of the password that was configured on the ISE as part of the device credentials.

Step 6 In the Confirm Password field, re-enter the password to confirm it.

Step 7 Click **Import**. The Identity By TrustSec pane reappears.

Step 8 Click **Apply** to save the changes.

The changes are saved to the running configuration.

Configuring the Security Exchange Protocol (SXP)

Configuring the Security Exchange Protocol (SXP) involves enabling the protocol in the ASA and setting the following default values for SXP:

- The source IP address of SXP connections
- The authentication password between SXP peers
- The retry interval for SXP connections
- The Cisco TrustSec SXP reconcile period

To configure the default settings for the ASA integration with Cisco TrustSec, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**. The Identity By TrustSec pane appears.

Step 2 Check the Enable Security Exchange Protocol checkbox to enable SXP. By default, SXP is disabled.

In multi-context mode, enabling SXP is done in the user context.

Step 3 In the Default Source field, enter the default local IP address for SXP connections. The IP address can be an IPv4 or IPv6 address.

Note The ASA determines the local IP address for an SXP connection as the outgoing interface IP address that is reachable by the peer IP address. If the configured local address is different from the outgoing interface IP address, the ASA cannot connect to the SXP peer and generates a system log message.

Step 4 In the Default Password field, enter the default password for TCP MD5 authentication with SXP peers. By default, SXP connections do not have a password set.

You can specify the password as an encrypted string up to 162 characters or an ASCII key string up to 80 characters. Configuring an encryption level for the password is optional. If you configure an encryption level, you can only set one level:

- Level 0—unencrypted cleartext
- Level 8—encrypted text

Step 5 In the Retry Timer field, enter the default time interval between ASA attempts to set up new SXP connections between SXP peers.

The ASA will continue to attempt to connect to new SXP peers until a successful connection is made. The retry timer is triggered as long as there is one SXP connection on the ASA that is not up.

Enter the retry timer value as a number of seconds in the range of 0 to 64000 seconds. If you specify 0 seconds, the timer never expires and the ASA will not attempt to connect to SXP peers. By default, the *timervalue* is 120 seconds.

When the retry timer expires, the ASA goes through the connection database and if the database contains any connections that are off or in a "pending on" state, the ASA restarts the retry timer.

Step 6 In the Reconcile Timer field, enter the default reconcile timer.

After an SXP peer terminates its SXP connection, the ASA starts a hold down timer. If an SXP peer connects while the hold down timer is running, the ASA starts the reconcile timer; then, the ASA updates the SXP mapping database to learn the latest mappings.

When the reconcile timer expires, the ASA scans the SXP mapping database to identify stale mapping entries (entries that were learned in a previous connection session). The ASA marks these connections as obsolete. When the reconcile timer expires, the ASA removes the obsolete entries from the SXP mapping database.

Enter the reconcile timer value as a number of seconds in the range of 1 to 64000 seconds. By default, the *timervalue* is 120 seconds.

Note You cannot specify 0 for the timer because specifying 0 would prevent the reconcile timer from starting. Not allowing the reconcile timer to run would keep stale entries for an undefined time and cause unexpected results from the policy enforcement.

Step 7 Click **Apply** to save the default settings.

The changes are saved to the running configuration.

Adding an SXP Connection Peer

SXP connections between peers are point-to-point and use TCP as the underlying transport protocol.

To add an SXP connection peer, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**. The Identity By TrustSec pane appears.

Step 2 If necessary, check the Enable Security Exchange Protocol checkbox to enable SXP.

Step 3 Click **Add**. The Add Connection dialog box appears.

Step 4 In the Peer IP Address field, enter the IPv4 or IPv6 address of the SXP peer. The peer IP address must be reachable from the ASA outgoing interface.

Step 5 **(Optional)** In the Source IP Address field, enter the local IPv4 or IPv6 address of the SXP connection. Specifying the source IP address is optional, however, specifying it safeguards misconfiguration.

Step 6 From the Password drop-down list, specify whether to use the authentication key for the SXP connection by selecting the following values:

- Default—Use the default password configured for SXP connections.

  See Configuring the Security Exchange Protocol (SXP), page 48-14.
- None—Do not use a password for the SXP connection.

Step 7 **(Optional)** From the Mode drop-down list, specify the mode of the SXP connection by selecting one of the following values:

- Local—Use the local SXP device.
- Peer—Use the peer SXP device.

Step 8 From the Role drop-down list, specify whether the ASA functions as a Speaker or Listener for the SXP connection:

- Speaker—The ASA can forward IP-SGT mappings to upstream devices.
- Listener—The ASA can receive IP-SGT mappings from downstream devices.

See About Speaker and Listener Roles on the ASA, page 48-6.

Step 9 Click **OK**. The peer appears in the Connection Peers list.

Step 10 Click **Apply** to save SXP peer settings.

The changes are saved to the running configuration.

Refreshing Environment Data

The ASA downloads environment data from the ISE, which includes the Security Group Tag (SGT) name table. The ASA automatically refreshes its environment data obtained from the ISE when you complete the following tasks on the ASA:

- Configure a AAA server to communicate with the ISE.
- Import a PAC file from the ISE.
- Identify the AAA server group that the ASA will use for retrieval of Cisco TrustSec environment data.

Normally, you will not need to manually refresh the environment data from the ISE; however, security groups can change on the ISE. These changes are not reflected on the ASA until you refresh the data in the ASA security group table. Refresh the data on the ASA to make sure any security group made on the ISE are reflected on the ASA.

Tip We recommend that you schedule policy configuration changes on the ISE and the manual data refresh on the ASA during a maintenance window. Handling policy configuration changes in this way maximizes the chances of security group names getting resolved and security policies becoming active immediately on the ASA.

Prerequisites

The ASA must be configured as a recognized Cisco TrustSec network device in the ISE and the ASA must have successfully imported a PAC file, so that the changes made for Cisco TrustSec are applied to the ASA.

Restrictions

- When the ASA is part of an HA configuration, you must refresh the environment data on the primary ASA device.
- When the ASA is part of a clustering configuration, you must refresh the environment data on the master device.

To Refresh Environment Data

In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**. The Identity By TrustSec pane appears. In the Server Group Setup area, click **Refresh Environment Data**.

The ASA refreshes the Cisco TrustSec environment data from the ISE and resets the reconcile timer to the configured default value.

Configuring the Security Policy

You can incorporate TrustSec policy in many ASA features. Any feature that uses extended ACLs (unless listed in this chapter as unsupported) can take advantage of TrustSec. You can now add security group arguments to extended ACLs, as well as traditional network-based parameters.

- To configure an extended ACL, see Chapter 19, "Adding an Extended Access Control List."
- To configure security group object groups, which can be used in the ACL, see the "Configuring Local User Groups" section on page 25-7.

For example, an access rule permits or denies traffic on an interface using network information. With TrustSec, you can now control access based on security group. See Chapter 51, "Configuring Access Rules." For example, you could create an access rule for sample_securitygroup1 10.0.0.0 255.0.0.0, meaning the security group could have any IP address on subnet 10.0.0.0/8.

You can configure security policies based on combinations of security group names (servers, users, unmanaged devices, etc.), user-based attributes, and traditional IP-address-based objects (IP address, Active Directory object, and FQDN). Security-group membership can extend beyond roles to include device and location attributes and is independent of user-group membership.

Monitoring the ASA Integrated with Cisco TrustSec

This section contains the following topics:

- Monitoring SXP Connections, page 48-17
- Monitoring Environment Data, page 48-18
- Monitoring Cisco TrustSec IP-SGT Mappings, page 48-19
- Monitoring the PAC File, page 48-20

Monitoring SXP Connections

To verify which SXP connections are up and running, you can view the Security Exchange Protocol monitoring panel in ASDM. This panel displays the SXP connections on the ASA for a particular user context when multi-context mode is used.

In the main ASDM application window, choose **Monitoring > Properties > Identity By TrustSec > Security Exchange Protocol**. The following monitoring panel appears.

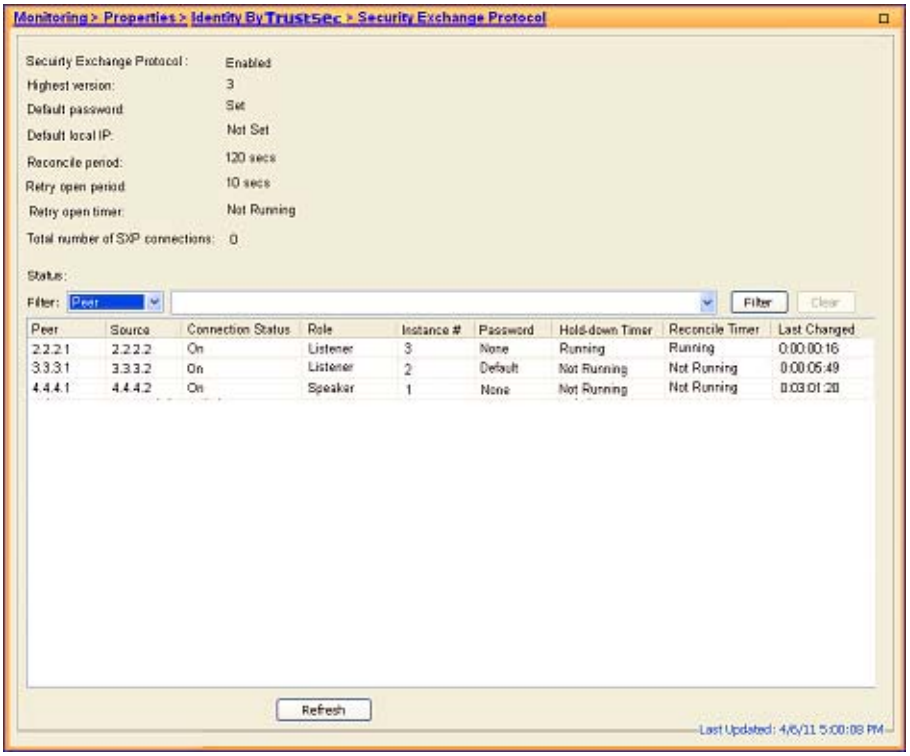

Alternatively, you can view the Connection monitoring panel in ASDM to display all connections configured for the ASA including the SXP connections. In the main ASDM application window, choose **Monitoring > Properties > Connections**. The following monitoring panel appears.

You can filter the IP-SGT mappings so that you view the data by SGT value, security group name, or IP address.

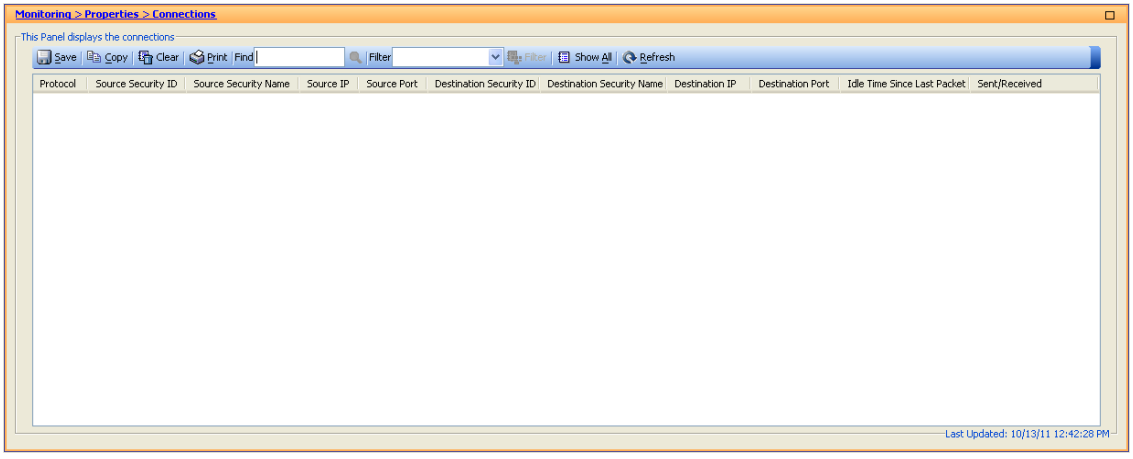

Monitoring Environment Data

To display the Cisco TrustSec environment information contained in security group table on the ASA, you can view the Environment Data monitoring panel in ASDM. The information in this panel includes the expiry timeout and security group name table. The security group table is populated with data from the ISE when you import the PAC file.

In the main ASDM application window, choose **Monitoring > Properties > Identity By TrustSec > Environment Data**. The following monitoring panel appears.

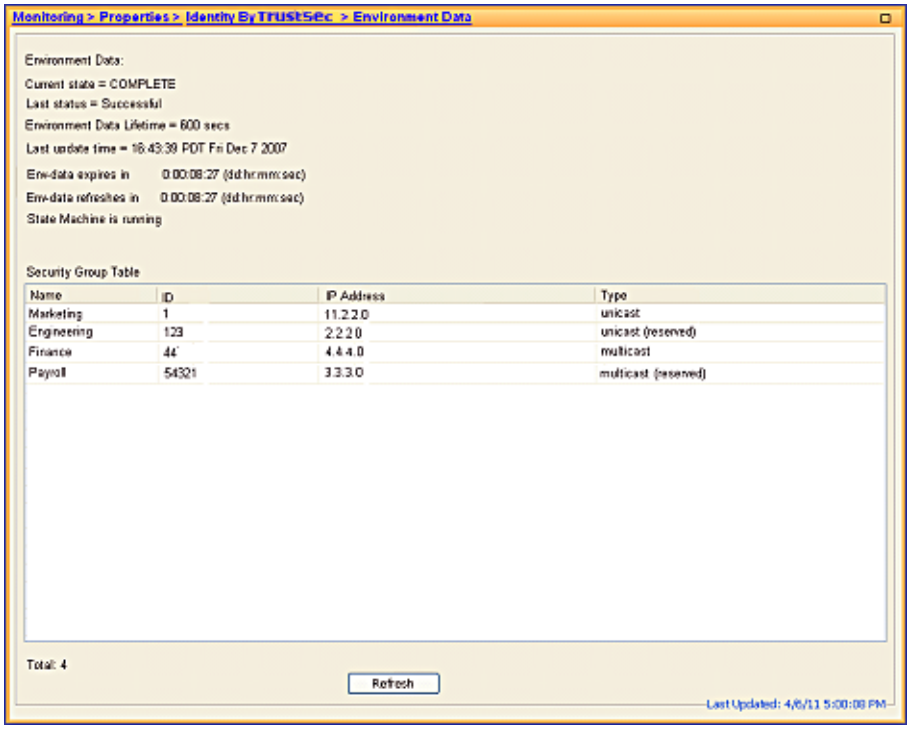

Monitoring Cisco TrustSec IP-SGT Mappings

To display the active IP-SGT mappings on the ASA consolidated from SXP, you can view the IP Mappings monitoring panel in ASDM.

In the main ASDM application window, choose **Monitoring > Properties > Identity By TrustSec > IP Mapping**. The following monitoring panel appears. You can filter the IP-SGT mappings so that you view the data by SGT value, security group name, or IP address.

Tip: Click **Where Used** to display where the selected security group object is used in an access list or nested in another security group object.

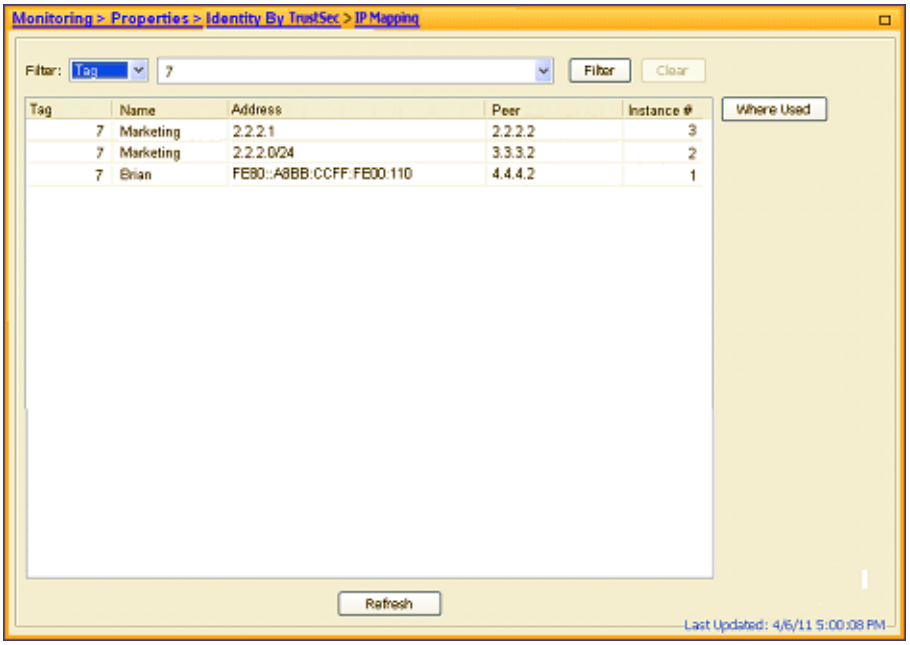

Monitoring the PAC File

To display the active PAC file imported into the ASA, you can view the PAC File monitoring panel in ASDM.

In the main ASDM application window, choose **Monitoring > Properties > Identity By TrustSec > PAC File**. The following monitoring panel appears.

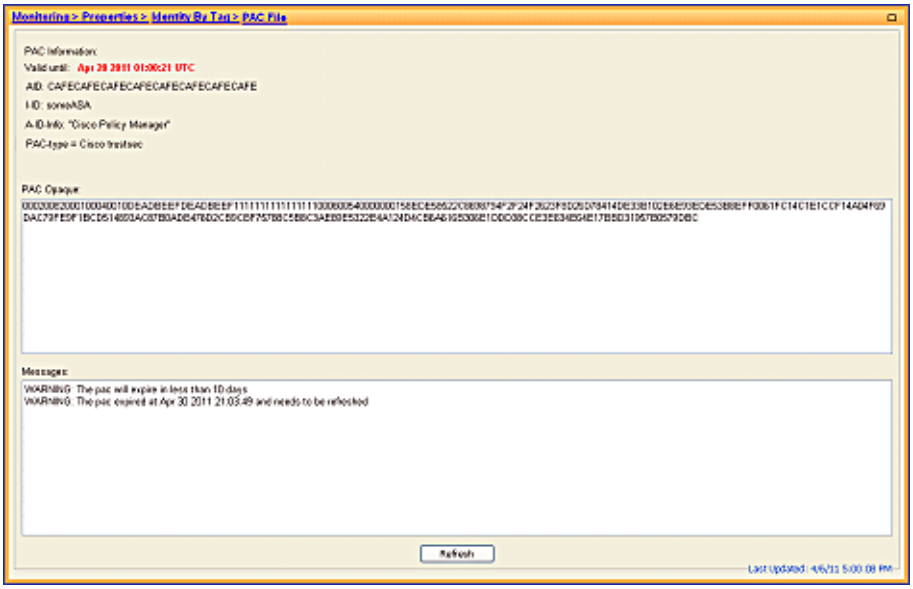

Feature History for the ASA-Cisco TrustSec Integration

Table 48-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 48-3 Feature History for the ASA-Cisco TrustSec Integration

Feature Name	Platform Releases	Feature Information
Cisco TrustSec Integration	9.0(1)	Cisco TrustSec provides an access-control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec solution, enforcement devices utilize a combination of user attributes and end-point attributes to make role-based and identity-based access control decisions. In this release, the ASA integrates with Cisco TrustSec to provide security group based policy enforcement. Access policies within the Cisco TrustSec domain are topology-independent, based on the roles of source and destination devices rather than on network IP addresses. The ASA can utilize the Cisco TrustSec solution for other types of security group based policies, such as application inspection; for example, you can configure a class map containing an access policy based on a security group. We introduced or modified the following screens: Configuration > Firewall > Identity By TrustSec Configuration > Firewall > Objects > Security Groups Object Groups Configuration > Firewall > Access Rules > Add Access Rules Monitoring > Properties > Identity By Tag

Configuring Digital Certificates

This chapter describes how to configure digital certificates and includes the following sections:

- Information About Digital Certificates, page 49-1
- Licensing Requirements for Digital Certificates, page 49-8
- Prerequisites for Local Certificates, page 49-8
- Guidelines and Limitations, page 49-9
- Configuring Digital Certificates, page 49-10
- Configuring CA Certificate Authentication, page 49-11
- Monitoring CRLs, page 49-18
- Configuring Identity Certificates Authentication, page 49-22
- Configuring Code Signer Certificates, page 49-27
- Authenticating Using the Local CA, page 49-29
- Managing the User Database, page 49-33
- Managing User Certificates, page 49-35
- Monitoring CRLs, page 49-36
- Feature History for Certificate Management, page 49-37

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 ASA. This configuration allows multiple identities, roots, and certificate hierarchies. The ASA evaluates third-party certificates against CRLs, also called authority revocation lists, all the way from the identity certificate up the chain of subordinate certificate authorities.

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 49-11.
- 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 49-22.
- *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 49-27.

The local CA integrates an independent certificate authority feature on the ASA, 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 49-29, the "Managing the User Database" section on page 49-33, and the "Managing User Certificates" section on page 49-35.

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.

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 ASA. 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 49-13.

CAs are responsible for managing certificate requests and issuing digital certificates. A digital certificate includes information that identifies a user or device, such as a name, serial number, company, department, or IP address. A digital certificate also includes a copy of the public key for the user or device. A CA can be a trusted third party, such as VeriSign, or a private (in-house) CA that you establish within your organization.

Tip For an example of a scenario that includes certificate configuration and load balancing, see the following URL: https://supportforums.cisco.com/docs/DOC-5964.

This section includes the following topics:

- Public Key Cryptography, page 49-3
- Certificate Scalability, page 49-3
- Key Pairs, page 49-4
- Trustpoints, page 49-4
- Revocation Checking, page 49-5

- The Local CA, page 49-7

Public Key Cryptography

Digital signatures, enabled by public key cryptography, provide a way to authenticate devices and users. In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other.

In simple terms, a signature is formed when data is encrypted with a private key. The signature is attached to the data and sent to the receiver. The receiver applies the public key of the sender to the data. If the signature sent with the data matches the result of applying the public key to the data, the validity of the message is established.

This process relies on the receiver having a copy of the public key of the sender and a high degree of certainty that this key belongs to the sender, not to someone pretending to be the sender.

Obtaining the public key of a sender is normally handled externally or through an operation performed at installation. For example, most web browsers are configured with the root certificates of several CAs by default. For VPN, the IKE protocol, a component of IPsec, can use digital signatures to authenticate peer devices before setting up security associations.

Certificate Scalability

Without digital certificates, you must manually configure each IPsec peer for each peer with which it communicates; as a result, each new peer that you add to a network would require a configuration change on each peer with which it needs to communicate securely.

When you use digital certificates, each peer is enrolled with a CA. When two peers try to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new peer is added to the network, you enroll that peer with a CA and none of the other peers need modification. When the new peer attempts an IPsec connection, certificates are automatically exchanged and the peer can be authenticated.

With a CA, a peer authenticates itself to the remote peer by sending a certificate to the remote peer and performing some public key cryptography. Each peer sends its unique certificate, which was issued by the CA. This process works because each certificate encapsulates the public key for the associated peer, each certificate is authenticated by the CA, and all participating peers recognize the CA as an authenticating authority. The process is called IKE with an RSA signature.

The peer can continue sending its certificate for multiple IPsec sessions, and to multiple IPsec peers, until the certificate expires. When its certificate expires, the peer administrator must obtain a new one from the CA.

CAs can also revoke certificates for peers that no longer participate in IPsec. Revoked certificates are not recognized as valid by other peers. Revoked certificates are listed in a CRL, which each peer may check before accepting a certificate from another peer.

Some CAs have an RA as part of their implementation. An RA is a server that acts as a proxy for the CA, so that CA functions can continue when the CA is unavailable.

Key Pairs

Key pairs are RSA keys, which have the following characteristics:

- RSA keys can be used for SSH or SSL.
- SCEP enrollment supports the certification of RSA keys.
- For the purposes of generating keys, the maximum key modulus for RSA keys is 2048 bits. The default size is 1024. Many SSL connections using identity certificates with RSA key pairs that exceed 1024 bits can cause a high CPU usage on the ASA and rejected clientless logins.
- For signature operations, the supported maximum key size is 4096 bits.
- You can generate a general purpose RSA key pair, used for both signing and encryption, or you can generate separate RSA key pairs for each purpose. Separate signing and encryption keys help to reduce exposure of the keys, because SSL uses a key for encryption but not signing. However, IKE uses a key for signing but not encryption. By using separate keys for each, exposure of the keys is minimized.

Trustpoints

Trustpoints let you manage and track CAs and certificates. A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one, enrolled identity certificate.

After you have defined a trustpoint, you can reference it by name in commands requiring that you specify a CA. You can configure many trustpoints.

Note If an ASA has multiple trustpoints that share the same CA, only one of these trustpoints sharing the CA can be used to validate user certificates. To control which trustpoint sharing a CA is used for validation of user certificates issued by that CA, use the **support-user-cert-validation** command.

For automatic enrollment, a trustpoint must be configured with an enrollment URL, and the CA that the trustpoint represents must be available on the network and must support SCEP.

You can export and import the keypair and issued certificates associated with a trustpoint in PKCS12 format. This format is useful to manually duplicate a trustpoint configuration on a different ASA.

Certificate Enrollment

The ASA needs a CA certificate for each trustpoint and one or two certificates for itself, depending upon the configuration of the keys used by the trustpoint. If the trustpoint uses separate RSA keys for signing and encryption, the ASA needs two certificates, one for each purpose. In other key configurations, only one certificate is needed.

The ASA supports automatic enrollment with SCEP and with manual enrollment, which lets you paste a base-64-encoded certificate directly into the terminal. For site-to-site VPNs, you must enroll each ASA. For remote access VPNs, you must enroll each ASA and each remote access VPN client.

Proxy for SCEP Requests

The ASA can proxy SCEP requests between AnyConnect and a third-party CA. The CA only needs to be accessible to the ASA if it is acting as the proxy. For the ASA to provide this service, the user must authenticate using any of the methods supported by AAA before the ASA sends an enrollment request. You can also use host scan and dynamic access policies to enforce rules of eligibility to enroll.

The ASA supports this feature only with an AnyConnect SSL or IKEv2 VPN session. It supports all SCEP-compliant CAs, including IOS CS, Windows Server 2003 CA, and Windows Server 2008 CA.

Clientless (browser-based) access does not support SCEP proxy, although WebLaunch—clientless-initiated AnyConnect—does support it.

The ASA does not support polling for certificates.

The ASA supports load balancing for this feature.

Revocation Checking

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 ASA to check that the CA has not revoked a certificate each time that it uses the certificate for authentication.

When you enable revocation checking, the ASA checks certificate revocation status during the PKI certificate validation process, which can use either CRL checking, OCSP, or both. OCSP is only used when the first method returns an error (for example, indicating that the server is unavailable).

With CRL checking, the ASA retrieves, parses, and caches CRLs, which provide a complete list of revoked (and unrevoked) certificates with their certificate serial numbers. The ASA evaluates certificates according to CRLs, also called authority revocation lists, from the identity certificate up the chain of subordinate certificate authorities.

OCSP offers a more scalable method of checking revocation status in that it localizes certificate status through a validation authority, which it queries for status of a specific certificate.

Supported CA Servers

The ASA supports the following CA servers:

Cisco IOS CS, ASA Local CA, and third-party X.509 compliant CA vendors including, but not limited to:

- Baltimore Technologies
- Entrust
- Digicert
- Geotrust
- GoDaddy
- iPlanet/Netscape
- Microsoft Certificate Services
- RSA Keon
- Thawte

- VeriSign

CRLs

CRLs provide the ASA with one way of determining whether a certificate that is within its valid time range has been revoked by the issuing CA. CRL configuration is part of configuration of a trustpoint.

You can configure the ASA to make CRL checks mandatory when authenticating a certificate by using the **revocation-check crl** command. You can also make the CRL check optional by using the **revocation-check crl none** command, which allows the certificate authentication to succeed when the CA is unavailable to provide updated CRL data.

The ASA can retrieve CRLs from CAs using HTTP, SCEP, or LDAP. CRLs retrieved for each trustpoint are cached for a configurable amount of time for each trustpoint.

When the ASA has cached a CRL for longer than the amount of time it is configured to cache CRLs, the ASA considers the CRL too old to be reliable, or "stale." The ASA tries to retrieve a newer version of the CRL the next time that a certificate authentication requires a check of the stale CRL.

The ASA caches CRLs for an amount of time determined by the following two factors:

- The number of minutes specified with the **cache-time** command. The default value is 60 minutes.
- The NextUpdate field in the CRLs retrieved, which may be absent from CRLs. You control whether the ASA requires and uses the NextUpdate field with the **enforcenextupdate** command.

The ASA uses these two factors in the following ways:

- If the NextUpdate field is not required, the ASA marks CRLs as stale after the length of time defined by the **cache-time** command.
- If the NextUpdate field is required, the ASA marks CRLs as stale at the sooner of the two times specified by the **cache-time** command and the NextUpdate field. For example, if the **cache-time** command is set to 100 minutes and the NextUpdate field specifies that the next update is 70 minutes away, the ASA marks CRLs as stale in 70 minutes.

If the ASA has insufficient memory to store all CRLs cached for a given trustpoint, it deletes the least recently used CRL to make room for a newly retrieved CRL.

OCSP

OCSP provides the ASA with a way of determining whether a certificate that is within its valid time range has been revoked by the issuing CA. OCSP configuration is part of trustpoint configuration.

OCSP localizes certificate status on a validation authority (an OCSP server, also called the *responder*) which the ASA queries for the status of a specific certificate. This method provides better scalability and more up-to-date revocation status than does CRL checking, and helps organizations with large PKI installations deploy and expand secure networks.

Note The ASA allows a five-second time skew for OCSP responses.

You can configure the ASA to make OCSP checks mandatory when authenticating a certificate by using the **revocation-check ocsp** command. You can also make the OCSP check optional by using the **revocation-check ocsp none** command, which allows the certificate authentication to succeed when the validation authority is unavailable to provide updated OCSP data.

OCSP provides three ways to define the OCSP server URL. The ASA uses these servers in the following order:

1. The OCSP URL defined in a match certificate override rule by using the **match certificate** command).
2. The OCSP URL configured by using the **ocsp url** command.
3. The AIA field of the client certificate.

Note To configure a trustpoint to validate a self-signed OCSP responder certificate, you import the self-signed responder certificate into its own trustpoint as a trusted CA certificate. Then you configure the **match certificate** command in the client certificate validating trustpoint to use the trustpoint that includes the self-signed OCSP responder certificate to validate the responder certificate. Use the same procedure for configuring validating responder certificates external to the validation path of the client certificate.

The OCSP server (responder) certificate usually signs the OCSP response. After receiving the response, the ASA tries to verify the responder certificate. The CA normally sets the lifetime of the OCSP responder certificate to a relatively short period to minimize the chance of being compromised. The CA usually also includes an ocsp-no-check extension in the responder certificate, which indicates that this certificate does not need revocation status checking. However, if this extension is not present, the ASA tries to check revocation status using the same method specified in the trustpoint. If the responder certificate is not verifiable, revocation checks fail. To avoid this possibility, use the **revocation-check none** command to configure the responder certificate validating trustpoint, and use the r**evocation-check ocsp** command to configure the client certificate.

The Local CA

The local CA performs the following tasks:

- Integrates basic certificate authority operation on the ASA.
- Deploys certificates.
- Provides secure revocation checking of issued certificates.
- Provides a certificate authority on the ASA for use with browser-based and client-based SSL VPN connections.
- Provides trusted digital certificates to users, without the need to rely on external certificate authorization.
- Provides a secure, in-house authority for certificate authentication and offers straightforward user enrollment by means of a website login.

Storage for Local CA Files

The ASA accesses and implements user information, issued certificates, and revocation lists using a local CA database. This database resides in local flash memory by default, or can be configured to reside on an external file system that is mounted and accessible to the ASA.

No limits exist on the number of users that can be stored in the local CA user database; however, if flash memory storage issues arise, syslogs are generated to alert the administrator to take action, and the local CA could be disabled until the storage issues are resolved. Flash memory can store a database with 3500 users or less; however, a database of more than 3500 users requires external storage.

The Local CA Server

After you configure a local CA server on the ASA, users can enroll for a certificate by logging into a website and entering a username and a one-time password that is provided by the local CA administrator to validate their eligibility for enrollment.

As shown in Figure 49-1, the local CA server resides on the ASA and handles enrollment requests from website users and CRL inquiries coming from other certificate validating devices and ASAs. Local CA database and configuration files are maintained either on the ASA flash memory (default storage) or on a separate storage device.

Figure 49-1 The Local CA

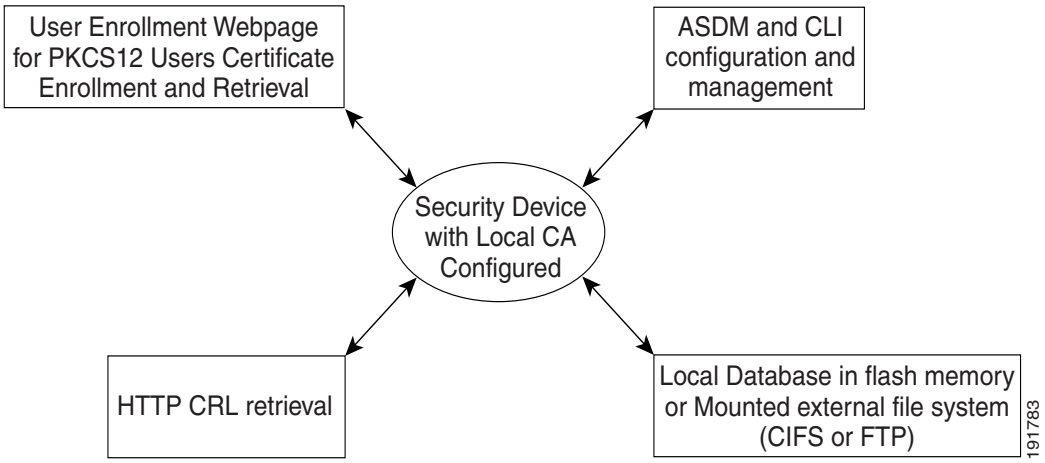

Licensing Requirements for Digital Certificates

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites for Local Certificates

Local certificates have the following prerequisites:

- Make sure that the ASA is configured correctly to support certificates. An incorrectly configured ASA can cause enrollment to fail or request a certificate that includes inaccurate information.
- Make sure that the hostname and domain name of the ASA are configured correctly. To view the currently configured hostname and domain name, enter the **show running-config** command. For information about configuring the hostname and domain name, see the "Configuring the Hostname, Domain Name, and Passwords" section on page 21-1.

- Make sure that the ASA clock is set accurately before configuring the CA. Certificates have a date and time that they become valid and expire. When the ASA enrolls with a CA and obtains a certificate, the ASA checks that the current time is within the valid range for the certificate. If it is outside that range, enrollment fails. For information about setting the clock, see the "Setting the Date and Time" section on page 21-2.

Prerequisites for SCEP Proxy Support

Configuring the ASA as a proxy to submit requests for third-party certificates has the following requirements:

- AnyConnect Secure Mobility Client 3.0 or later must be running at the endpoint.
- The authentication method, configured in the connection profile for your group policy, must be set to use both AAA and certificate authentication.
- An SSL port must be open for IKEv2 VPN connections.
- The CA must be in auto-grant mode.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- Supported in single and multiple context mode for a local CA.
- Supported in single context mode only for third-party CAs.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

- Does not support replicating sessions in Stateful Failover.
- Does not support Active/Active or Active/Standby failover.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

- For ASAs that are configured as CA servers or clients, limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038. This guideline also applies to imported certificates from third-party vendors.
- You cannot configure the local CA when failover is enabled. You can only configure the local CA server for standalone ASAs without failover. For more information, see CSCty43366.
- When a certificate enrollment is completed, the ASA stores a PKCS12 file containing the user's keypair and certificate chain, which requires about 2 KB of flash memory or disk space per enrollment. The actual amount of disk space depends on the configured RSA key size and certificate

fields. Keep this guideline in mind when adding a large number of pending certificate enrollments on an ASA with a limited amount of available flash memory, because these PKCS12 files are stored in flash memory for the duration of the configured enrollment retrieval timeout.

- The **lifetime ca-certificate** command takes effect when the local CA server certificate is first generated (that is, when you initially configure the local CA server and issue the **no shutdown** command). When the CA certificate expires, the configured lifetime value is used to generate the new CA certificate. You cannot change the lifetime value for existing CA certificates.
- You should configure the ASA to use an identity certificate to protect ASDM traffic and HTTPS traffic to the management interface. Identity certificates that are automatically generated with SCEP are regenerated after each reboot, so make sure that you manually install your own identity certificates. For an example of this procedure that applies only to SSL, see the following URL: http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a00809fcf91.shtml.
- The ASA and the AnyConnect clients can only validate certificates in which the X520Serialnumber field (the serial number in the Subject Name) is in PrintableString format. If the serial number format uses encoding such as UTF8, the certificate authorization will fail.

Configuring Digital Certificates

This section describes how to configure local CA certificates. Make sure that you follow the sequence of tasks listed to correctly configure this type of digital certificate. This section includes the following topics:

- Configuring CA Certificate Authentication, page 49-11
- Configuring CA Certificates for Revocation, page 49-19
- Configuring CRL Retrieval Policy, page 49-19
- Configuring CRL Retrieval Methods, page 49-20
- Configuring OCSP Rules, page 49-20
- Configuring Advanced CRL and OCSP Settings, page 49-21

This section describes how to configure digital certificates for the ASA Services Module and includes the following topics:

- Configuring CA Certificate Authentication, page 49-13
- Adding or Installing a CA Certificate, page 49-13
- Editing or Removing a CA Certificate Configuration, page 49-14
- Showing CA Certificate Details, page 49-14
- Configuring CA Certificates for Revocation, page 49-15
- Configuring CRL Retrieval Policy, page 49-15
- Configuring CRL Retrieval Methods, page 49-16
- Configuring OCSP Rules, page 49-16
- Configuring Advanced CRL and OCSP Settings, page 49-17

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 ASA.
- Create a new certificate configuration.
- Edit an existing certificate configuration.
- Obtain a CA certificate manually and import it.
- Have the ASA 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, page 49-11
- Editing or Removing a CA Certificate Configuration, page 49-12
- Showing CA Certificate Details, page 49-13

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 The Certificate Installation dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.

Step 6 To enroll manually, click the **Paste certificate in PEM format** radio button.

Step 7 Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided, then click **Install Certificate**.

Step 8 The Certificate Installation dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.

Step 9 To enroll automatically, click the **Use SCEP** radio button. The ASA 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.

Note See Prerequisites for SCEP Proxy Support when choosing to use the SCEP method to install certficates.

Step 10 To display additional configuration options for new and existing certificates, click **More Options**.

The Configuration Options for CA Certificates pane appears.

Step 11 To continue, see the "Editing or Removing a CA Certificate Configuration" section on page 49-12.

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 CRL Retrieval Policy" section on page 49-19
- "Configuring CRL Retrieval Methods" section on page 49-20
- "Configuring OCSP Rules" section on page 49-20
- "Configuring Advanced CRL and OCSP Settings" section on page 49-21

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.

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 ASA.
- Create a new certificate configuration.
- Edit an existing certificate configuration.
- Obtain a CA certificate manually and import it.
- 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.

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.

To add or install a CA certificate, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Device Management > 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 display additional configuration options for new and existing certificates, click **More Options**.

The Configuration Options for CA Certificates pane appears.

Step 8 Make your selections, and then click OK. To continue, see the "Editing or Removing a CA Certificate Configuration" section on page 49-14.

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 CRL Retrieval Policy" section on page 49-15
- "Configuring CRL Retrieval Methods" section on page 49-16
- "Configuring OCSP Rules" section on page 49-16
- "Configuring Advanced CRL and OCSP Settings" section on page 49-17

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.

Configuring CA Certificates for Revocation

To configure CA certificates for revocation, perform the following site-to-site task in either single or muliple context mode:

Step 1 In the ASDM application window, choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the Install Certificates dialog box. Then click **More Options**.

Step 2 In the Configuration Options for CA Certificates pane, click the **Revocation Check** tab.

Step 3 To disable revocation checking of certificates, click the **Do not check certificates for revocation** radio button.

Step 4 To select one or more revocation checking methods (CRL or OCSP), click the **Check certificates for revocation** radio button.

Step 5 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 6 Check the **Consider certificate valid if revocation checking returns errors** check box to ignore revocation checking errors during certificate validation.

Step 7 Click **OK** to close the Revocation Check tab. Alternatively, to continue, see the "Configuring CRL Retrieval Policy" section on page 49-15.

Configuring CRL Retrieval Policy

To configure the CRL retrieval policy, perform the following steps:

Step 1 In the ASDM application window, choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the Install Certificates dialog box. Then click **More Options**.

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 49-16.

Configuring CRL Retrieval Methods

To configure CRL retrieval methods, perform the following steps:

Step 1 In the ASDM application window, choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the Install Certificates dialog box. Then click **More Options**.

Step 2 In the Configuration Options for CA Certificates pane, click the **CRL Retrieval Methods** tab.

Step 3 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.

Step 4 Click **OK** to close this tab. Alternatively, to continue, see the "Configuring OCSP Rules" section on page 49-16.

Configuring OCSP Rules

The ASA 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 > Site-to-Site VPN > Advanced > 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 ASDM application window, choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the Install Certificates dialog box. Then click **More Options**.

Step 2 In the Configuration Options for CA Certificates pane, click the **OCSP Rules** tab.

Step 3 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 ASA 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 4 To add a new OCSP rule, click **Add**.

The Add OCSP Rule dialog box appears.

Step 5 Choose the certificate map to use from the drop-down list.

Step 6 Choose the certificate to use from the drop-down list.

Step 7 Enter the priority number for the rule.

Step 8 Enter the URL of the OCSP server for this certificate.

Step 9 When you are done, click **OK** to close this dialog box.

The newly added OCSP rule appears in the list.

Step 10 To edit an existing OCSP rule, select it, and then click **Edit**.

Step 11 To delete an OCSP rule, select it, and then click **Delete**.

Step 12 Click **OK** to close this tab. Alternatively, to continue, see the "Configuring Advanced CRL and OCSP Settings" section on page 49-17.

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 ASA to check that the CA has not revoked the certificate being verified. The ASA 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 ASDM application window, choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the Install Certificates dialog box. Then click **More Options**.

Step 2 In the Configuration Options for CA Certificates pane, click the **Advanced** tab.

Step 3 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 ASA 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 ASA removes the least recently used CRL until more space becomes available.

Step 4 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 5 In the OCSP Options area, enter the URL for the OCSP server. The ASA 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 user certificate

Step 6 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 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 ASA 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 ASA 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 "Monitoring CRLs" section on page 49-18.

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.

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 that 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 49-19.

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 49-20.

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 49-20.

Configuring OCSP Rules

The ASA 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 ASA 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 49-21.

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 ASA to check that the CA has not revoked the certificate being verified. The ASA 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 ASA 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 ASA 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 ASA 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 buttons 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 ASA 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 ASA 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 49-22.

Configuring Identity Certificates Authentication

An identity certificate can be used to authenticate VPN access through the ASA. 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 49-22
- Showing Identity Certificate Details, page 49-24
- Deleting an Identity Certificate, page 49-25
- Exporting an Identity Certificate, page 49-25
- Generating a Certificate Signing Request, page 49-25
- Installing Identity Certificates, page 49-26

Adding or Importing an Identity Certificate

To add or import a new identity certificate configuration, perform the following:

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 (PKCS12 format with Certificate(s) + Private Key)** 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 Choose the **RSA** or **ECDSA** key type.

Note The 4096-bit RSA keys are only supported on the 5580, 5585, or later platforms. If you are configuring a cryptography map to use an RSA or ECDSA trustpoint for authentication, you must first generate the key set. You can then create the trustpoint and reference it in the tunnel group configuration.

Step 9 To use the default key pair name, click the **Use default keypair name** radio button.

Step 10 To use a new key pair name, click the **Enter a new key pair name** radio button, and type the new name. The ASA supports multiple key pairs.

Step 11 Choose the modulus size from the drop-down list. If you are notsure of the modulus size, consult Entrust.

Step 12 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 ASA 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 13 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 14 Click **OK** when you are done to close the Key Pair Details dialog box.

Step 15 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 16 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 17 Click **OK** when you are done to close the Certificate Subject DN dialog box.

Step 18 To create self-signed certificates, check the **Generate self-signed certificate** check box.

Step 19 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 20 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 21 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 ASA IP address on the network in four-part, dotted-decimal notation.
- To add the ASA serial number to the certificate parameters, check the **Include serial number of the device** check box.

Step 22 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 23 Click the **SCEP Challenge Password** tab, and then enter the following information:

- The SCEP password
- The SCEP password confirmation

Step 24 Click **OK** when you are done to close the Advanced Options dialog box.

Step 25 Click **Add Certificate** in the Add Identity Certificate dialog box.

The new identity certificate appears in the Identity Certificates list.

Step 26 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 ASAs; 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 ASAs.

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

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 49-22. When you generate the key pair, you can send it to the ASA 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 ASA.

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 ASA 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 ASA. 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 49-22.

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 ASA configuration.

What to Do Next

See the "Configuring Code Signer Certificates" section on page 49-27.

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 49-28
- Deleting a Code Signer Certificate, page 49-28
- Importing a Code Signer Certificate, page 49-28
- Exporting a Code Signer Certificate, page 49-29

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 49-29.

Authenticating Using the Local CA

The local CA provides a secure, configurable in-house authority that resides on the ASA 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 ASA, 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 49-30

- Deleting the Local CA Server, page 49-32

Configuring the Local CA Server

To configure a local CA server on the ASA, perform the following steps:

Step 1 Choose **Configuration > Remote Access VPN > Certificate Management > Local Certificate Authority > CA Server**.

Step 2 To activate the local CA server, check the **Enable Certificate Authority Server** check box. The default setting is disabled (unchecked). After you enable the local CA server, the ASA generates the local CA server certificate, key pair, and necessary database files, 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 3 When you enable the local CA for the first time, you must enter and confirm 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 4 Click **Apply** to save the local CA certificate and key pair, so the configuration is not lost if you reboot the ASA.

Step 5 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 ASA by unchecking the **Enable Certificate Authority Server** check box. 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 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 6 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 7 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). Make sure that you limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038.

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 Latest ASDM Syslog Messages pane:

```
%ASA-1-717049: Local CA Server certificate is due to expire in 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 8 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). Make sure that you limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038.

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 the format "adminname@hostname.com." 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 9 To configure additional options, click the **More Options** drop-down arrow.

Step 10 Enter the CRL distribution point, which is the CRL location on the ASA. The default location is http://hostname.domain/+CSCOCA+/asa_ca.crl.

Step 11 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 ASA matches the path, /user8/my_crl_file to the configured URL. When the path matches, the ASA returns the stored CRL file.

Step 12 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 13 Enter the database storage location to specify a storage area for the local CA configuration and data files. The ASA 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 14 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 15 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 16 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 ASA is through a web browser that uses standard HTML.

Step 17 Enter the length of time that a one-time password e-mailed to an enrolling user is valid. The default is 72 hours.

Step 18 Enter the number of days before expiration reminders are e-mailed to users. The default is 14 days.

Step 19 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 ASA, 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 49-33.

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 49-33
- Sending an Initial OTP or Replacing OTPs, page 49-34
- Editing a Local CA User, page 49-34
- Deleting a Local CA User, page 49-35
- Allowing User Enrollment, page 49-35
- Viewing or Regenerating an OTP, page 49-35

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 49-35.

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 the user certificate lifetime period runs out, to remove user 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 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 49-36.

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 2010
    NextUpdate: 15:58:34 UTC Nov 11 2010
    Cached Until: 15:58:34 UTC Nov 11 2010
    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 49-1 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. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 49-1 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. We introduced the following screens: Configuration > Remote Access VPN > Certificate Management Configuration > Site-to-Site VPN > Certificate Management. We introduced or modified the following screens: Configuration > Firewall > Advanced > Certificate Management > CA Certificates Configuration > Device Management > Certificate Management > CA Certificates.
SCEP proxy	8.4(1)	We introduced this feature, which provides secure deployment of device certificates from third-party CAs.

PART 9

Configuring Access Control

Configuring Access Rules

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

- Information About Access Rules, page 51-1
- Licensing Requirements for Access Rules, page 51-6
- Guidelines and Limitations, page 51-6
- Default Settings, page 51-7
- Configuring Access Rules, page 51-7
- Feature History for Access Rules, page 51-13

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

To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to Chapter 52, "Configuring Management Access."

Information About Access Rules

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

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

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

This section includes the following topics:

- General Information About Rules, page 51-2
- Information About Access Rules, page 51-4
- Information About EtherType Rules, page 51-5

General Information About Rules

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

- Implicit Permits, page 51-2
- Information About Interface Access Rules and Global Access Rules, page 51-2
- Using Access Rules and EtherType Rules on the Same Interface, page 51-2
- Rule Order, page 51-3
- Implicit Deny, page 51-3
- Using Remarks, page 51-3
- NAT and Access Rules, page 51-3
- Inbound and Outbound Rules, page 51-3
- Information About Access Rules, page 51-4

Implicit Permits

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

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

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

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

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

- BPDUs in both directions.

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

Information About Interface Access Rules and Global Access Rules

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

Note Global access rules apply only to inbound traffic. See the "Inbound and Outbound Rules" section on page 51-3.

Using Access Rules and EtherType Rules on the Same Interface

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

Rule Order

The order of rules is important. When the ASA decides whether to forward or drop a packet, the ASA tests the packet against each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create an access rule at the beginning that explicitly permits all traffic for an interface, no further rules are ever checked. For more information, see the "Implicit Deny" section on page 51-3.

You can disable a rule by making it inactive.

Implicit Deny

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

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

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

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

Using Remarks

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

NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

Inbound and Outbound Rules

The ASA supports two types of access lists:

- Inbound—Inbound access rules apply to traffic as it enters an interface. Global access rules are always inbound.
- Outbound—Outbound access lists apply to traffic as it exits an interface.

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

An outbound access list 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 access lists to restrict access, you can create a single outbound access list that allows only the specified hosts. (See Figure 51-1.) The outbound access list prevents any other hosts from reaching the outside network.

Figure 51-1 Outbound Access List

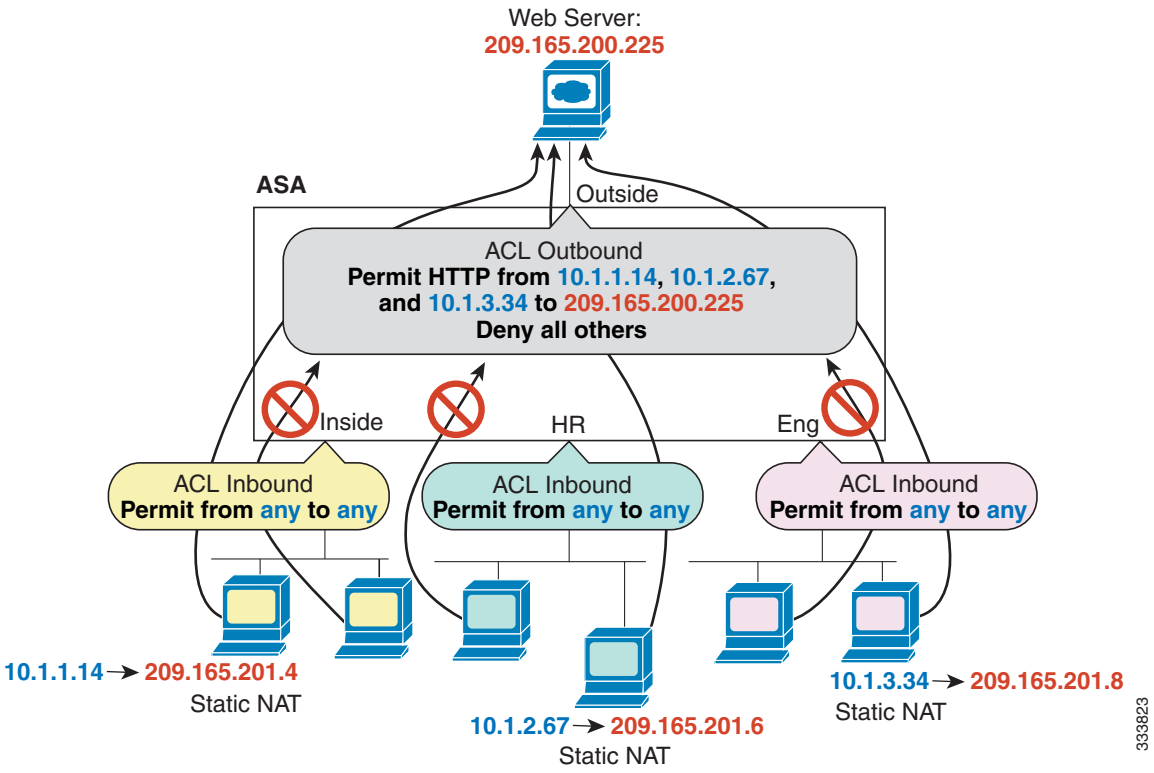

Information About Access Rules

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

- Access Rules for Returning Traffic, page 51-5
- Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules, page 51-5
- Management Access Rules, page 51-5

Access Rules for Returning Traffic

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

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

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

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Transparent firewall mode can allow any IP traffic through. This feature is especially useful in multiple context mode, which does not allow dynamic routing, for example.

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

Table 51-1 lists common traffic types that you can allow through the transparent firewall.

Table 51-1 Transparent Firewall Special Traffic

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

Management Access Rules

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

Information About EtherType Rules

This section describes EtherType rules and includes the following topics:

- Supported EtherTypes and Other Traffic, page 51-6
- Access Rules for Returning Traffic, page 51-6
- Allowing MPLS, page 51-6

Supported EtherTypes and Other Traffic

An EtherType rule controls the following:

- EtherType identified by a 16-bit hexadecimal number, including common types IPX and MPLS unicast or multicast.
- Ethernet V2 frames.
- BPDUs, which are permitted by default. BPDUs are SNAP-encapsulated, and the ASA is designed to specifically handle BPDUs.
- Trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the ASA modifies the payload with the outgoing VLAN if you allow BPDUs.
- IS-IS (supported in Version 8.4(5) only).

The following types of traffic are not supported:

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

Access Rules for Returning Traffic

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

Allowing MPLS

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

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

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

Or

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

Licensing Requirements for Access Rules

Model	License Requirement
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 modes.

IPv6 Guidelines

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

Default Settings

See the "Implicit Permits" section on page 51-2.

Configuring Access Rules

This section includes the following topics:

- Adding an Access Rule, page 51-7
- Adding an EtherType Rule (Transparent Mode Only), page 51-8
- Configuring Management Access Rules, page 51-9
- Advanced Access Rule Configuration, page 51-10
- Configuring HTTP Redirect, page 51-12

Adding an Access Rule

To apply an access rule, perform the following steps.

Detailed Steps

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

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

The Add Access Rule dialog box appears.

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

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

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

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

For more information about enabling IPv6 on an interface, see the "Configuring IPv6 Addressing" section on page 17-15.

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

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

See the "Configuring Identity-Based Security Policy" section on page 47-15 for more information.

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

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

Step 9 Select the service type.

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

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

   The Browse Time Range dialog box appears.

b. Click **Add**.

   The Add Time Range dialog box appears.

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

d. Choose the Start Time and the End Time.

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

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

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

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

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

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

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

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

Adding an EtherType Rule (Transparent Mode Only)

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

For more information about EtherType rules, see the "Information About Access Rules" section on page 51-1.

To add an EtherType rule, perform the following steps:

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

Step 2 Click **Add**.

The Add EtherType rules window appears.

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

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

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

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

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

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

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

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

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

Step 9 Click **OK**.

Configuring Management Access Rules

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

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

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

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

The Add Management Access Rule dialog box appears.

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

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

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

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

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

Note IPv6 must be enabled on at least one interface before you can configure an extended ACL with an IPv6 address. For more information about enabling IPv6 on an interface, see the "Configuring IPv6 Addressing" section on page 17-15

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

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

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

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

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

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

  The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.
- To configure the logging interval (if you enable logging and choose a non-default setting), enter a value in seconds in the Logging Interval field.
- To select a predefined time range for this rule, from the Time Range drop-down list, choose a time range; or click the ellipsis (...) to browse for a time range. You can also specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active.

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

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

Advanced Access Rule Configuration

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

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

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

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

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

Prerequisites

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

Fields

- Maximum Deny-flows—The maximum number of deny flows permitted before the ASA stops logging, between 1 and the default value. The default is 4096.
- Alert Interval—The amount of time (1-3600 seconds) between system log messages (number 106101) that identify that the maximum number of deny flows was reached. The default is 300 seconds.
- Per User Override table—Specifies the state of the per user override feature. If the per user override feature is enabled on the inbound access rule, the access rule provided by a RADIUS server replaces the access rule configured on that interface. If the per user override feature is disabled, the access rule provided by the RADIUS server is combined with the access rule configured on that interface. If the inbound access rule is not configured for the interface, per user override cannot be configured.
-

  For VPN remote access traffic, the behavior depends on whether there is a VPN filter applied in the group policy (see the Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > General > More Options > Filter field) and whether you set the Per User Override option:

  – No Per User Override, no VPN filter —Traffic is matched against the interface ACL (per the default **Enable inbound VPN sessions to bypass interface access lists** setting (disabled) on the Configurartion > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles pane).

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

  – Per User Override, VPN filter —Traffic is matched against the VPN filter only.
- Object Group Search Setting—Reduces the amount of memory used to store service rules, but lengthens the amount of time to search for a matching access rule.

Access Rule Explosion

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

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

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

Step 2 Click the **Advanced** button.

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

Configuring HTTP Redirect

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

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

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

Edit HTTP/HTTPS Settings

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

Fields

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

- Interface—Identifies the interface on which the ASA redirects or does not redirect HTTP requests to HTTPS.
- Redirect HTTP to HTTPS—Check to redirect HTTP requests to HTTPS, or uncheck to not redirect HTTP requests to HTTPS.
- HTTP Port—Identifies the port from which the interface redirects HTTP connections. By default it listens to port 80.

For more information about access rules, see the "Information About Access Rules" section on page 51-1.

Feature History for Access Rules

Table 51-2 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 51-2 Feature History for Access Rules

Feature Name	Platform Releases	Feature Information
Interface access rules	7.0(1)	Controlling network access through the ASA using access lists. We introduced the following screen: Configuration > Firewall > Access Rules.
Global access rules	8.3(1)	Global access rules were introduced. We modified the following screen: Configuration > Firewall > Access Rules.
Support for Identity Firewall	8.4(2)	You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication.
EtherType ACL support for IS-IS traffic (transparent firewall mode)	8.4(5)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules. *Not available in Version 8.5(1), 8.6(1), or 9.0(1).*
Support for TrustSec	9.0(1)	You can now use TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules.

Table 51-2 Feature History for Access Rules (continued)

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

Configuring Management Access

This chapter describes how to access the ASA for system management through Telnet, SSH, and HTTPS (using ASDM), how to authenticate and authorize users, and how to create login banners.

This chapter includes the following sections:

- Configuring ASA Access for ASDM, Telnet, or SSH, page 52-1
- Configuring CLI Parameters, page 52-5
- Configuring File Access, page 52-8
- Configuring ICMP Access, page 52-12
- Configuring Management Access Over a VPN Tunnel, page 52-14
- Configuring AAA for System Administrators, page 52-15
- Monitoring Device Access, page 52-31
- Feature History for Management Access, page 52-32

Note To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to the sections in this chapter.

Configuring ASA Access for ASDM, Telnet, or SSH

This section describes how to allow clients to access the ASA using ASDM, Telnet, or SSH and includes the following topics:

- Licensing Requirements for ASA Access for ASDM, Telnet, or SSH, page 52-1
- Guidelines and Limitations, page 52-2
- Configuring Management Access, page 52-3
- Using a Telnet Client, page 52-4
- Using an SSH Client, page 52-4

Licensing Requirements for ASA Access for ASDM, Telnet, or SSH

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.

IPv6 Guidelines

Supports IPv6.

Model Guidelines

For the ASASM, a session from the switch to the ASASM is a Telnet session, but Telnet access configuration according to this section is not required.

Additional Guidelines

- You cannot use Telnet to the lowest security interface unless you use Telnet inside a VPN tunnel.
- Management access to an interface other than the one from which you entered the ASA is not supported. For example, if your management host is located on the outside interface, you can only initiate a management connection directly to the outside interface. The only exception to this rule is through a VPN connection. See the "Configuring Management Access Over a VPN Tunnel" section on page 52-14.
- The ASA allows:
  - A maximum of 5 concurrent Telnet connections per context, if available, with a maximum of 100 connections divided among all contexts.
  - A maximum of 5 concurrent SSH connections per context, if available, with a maximum of 100 connections divided among all contexts.
  - A maximum of 5 concurrent ASDM instances per context, if available, with a maximum of 32 ASDM instances among all contexts.
- The ASA supports the SSH remote shell functionality provided in SSH Versions 1 and 2 and supports DES and 3DES ciphers.
- XML management over SSL and SSH is not supported.
- (8.4 and later) The SSH default username is no longer supported. You can no longer connect to the ASA using SSH with the **pix** or **asa** username and the login password. To use SSH, you must configure AAA authentication using Configuration > Device Management > Users/AAA > AAA Access > Authentication; then define a local user by choosing Configuration > Device Management > Users/AAA > User Accounts. If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.

- If you cannot make a Telnet or SSH connection to the ASA interface, make sure that you have enabled Telnet or SSH to the ASA according to the instructions in the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 52-1.

Configuring Management Access

To identify the client IP addresses allowed to connect to the ASA using Telnet, SSH, or ASDM, perform the following steps:

Detailed Steps

Step 1 Choose **Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH**, and click **Add**.

The Add Device Access Configuration dialog box appears.

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, enter 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 Click **OK**.

Step 7 Configure HTTP Settings.

a. Enable HTTP Server—Enable the HTTP server for ASDM access. This is enabled by default.

b. (Optional) Port Number—The default port is 443.

c. (Optional) Idle Timeout—The default idle timeout is 20 minutes.

d. (Optional) Session Timeout—By default, the session timeout is disabled. ASDM connections have no session time limit.

e. (Optional) Require client certificate to access ASDM on the following interfaces—Spercify the interface from the drop-down list.

Step 8 (Optional) Configure Telnet Settings.

a. Telnet Timeout—The default timeout value is 5 minutes.

Step 9 (Optional) Configure SSH Settings.

a. Allowed SSH Version(s)—The default value is 1 & 2.

b. SSH Timeout—The default timeout value is 5 minutes.

Step 10 Click **Apply**.

The changes are saved to the running configuration.

Step 11 (Required for SSH) You must also configure SSH authentication.

a. Choose **Configuration > Device Management > Users/AAA > AAA Access > Authentication**.

b. Check the **SSH** check box.

c. From the Server Group drop-down list, choose an already configured AAA server group name or the **LOCAL** database. To add AAA server groups, see the "Configuring AAA Server Groups" section on page 46-11.

d. (Optional) If you chose a AAA server group, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Check 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 ASA prompt does not give any indication of which method is being used.

e. Click **Apply**.

f. If you chose the LOCAL database, add a local user. Choose **Configuration > Device Management > Users/AAA > User Accounts**, and then click **Add**.

The Add User Account-Identity dialog box appears.

g. In the Username field, enter a username from 4 to 64 characters long.

h. In the Password field, enter a password between 3 and 32 characters. Passwords are case-sensitive.

i. In the Confirm Password field, reenter the password.

For information about other fields, see the "Adding a User Account to the Local Database" section on page 46-20.

j. Click **OK,** then click **Apply**.

Using a Telnet Client

To gain access to the ASA CLI using Telnet, enter the login password. If you configure Telnet authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20), then enter the username and password defined by the AAA server or local database.

Using an SSH Client

In the SSH client on your management host, enter the username and password. When starting an SSH session, a dot (.) displays on the ASA console before the following SSH user authentication prompt appears:

```
hostname(config)#.
```

The display of the dot does not affect the functionality of SSH. The dot appears at the console when generating a server key or decrypting a message using private keys during SSH key exchange before user authentication occurs. These tasks can take up to two minutes or longer. The dot is a progress indicator that verifies that the ASA is busy and has not hung.

Configuring CLI Parameters

This section includes the following topics:

- Licensing Requirements for CLI Parameters, page 52-5
- Guidelines and Limitations, page 52-5
- Configuring a Login Banner, page 52-5
- Customizing a CLI Prompt, page 52-6
- Changing the Console Timeout, page 52-8

Licensing Requirements for CLI Parameters

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.

Configuring a Login Banner

You can configure a message to display when a user connects to the ASA, before a user logs in, or before a user enters privileged EXEC mode.

Restrictions

After a banner is added, Telnet or SSH sessions to ASA may close if:

- There is not enough system memory available to process the banner message(s).
- A TCP write error occurs when trying to display banner message(s).

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.

To configure a login banner, perform the following steps:

Detailed Steps

Step 1 Choose **Configuration > Device Management > Management Access > Command Line (CLI) > Banner**, then add your banner text to the field for the type of banner that you are creating for the CLI:

- The session (exec) banner appears when a user accesses privileged EXEC mode at the CLI.
- The login banner appears when a user logs in to the CLI.
- The message-of-the-day (motd) banner appears when a user first connects to the CLI.
- The ASDM banner appears when a user connects to ASDM, after user authentication. The user is given two options for dismissing the banner:
  - Continue—Dismiss the banner and complete login.
  - Disconnect—Dismiss the banner and terminate the connection.
- Only ASCII characters are allowed, including a 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 ASA 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.

Step 2 Click **Apply**.

The new banner is 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 ASA. In multiple context mode, the prompt also displays the context name. You can display the following items in the CLI prompt:

cluster-unit	(Single and multiple mode) Displays the cluster unit name. Each unit in a cluster can have a unique name.
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 appear 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 another inactive 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 condition might occur when there is an interface failure above the threshold on the standby unit. Shows the role (master or slave) of a unit in a cluster. For example, in the prompt ciscoasa/cl2/slave, the hostname is ciscoasa, the unit name is cl2, and the state name is slave.

Detailed Steps

To customize the CLI prompt, perform the following steps:

Step 1 Choose **Configuration > Device Management > Management Access > Command Line (CLI) > CLI Prompt**, then 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

The console timeout sets how long a connection can remain in privileged EXEC mode or configuration mode; when the timeout is reached, the session drops into user EXEC mode. By default, the session does not time out. This setting does not affect how long you can remain connected to the console port, which never times out.

To change the console timeout, perform the following steps:

Detailed Steps

Step 1 To define a new timeout value in minutes, choose **Configuration > Device Management > Management Access > Command Line (CLI) > Console Timeout**.

Step 2 To specify an unlimited amount of time, enter **0**. The default value is 0.

Step 3 Click **Apply**.

The timeout value is changed and the change is saved to the running configuration.

Configuring File Access

This section includes the following topics:

- Licensing Requirements for File Access, page 52-8
- Guidelines and Limitations, page 52-8
- Configuring the FTP Client Mode, page 52-9
- Configuring the ASA as a Secure Copy Server, page 52-9
- Configuring the ASA as a TFTP Client, page 52-10
- Adding Mount Points, page 52-10

Licensing Requirements for File Access

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.

Configuring the FTP Client Mode

The ASA 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 the **Specify FTP mode as passive** check box.

Step 2 Click **Apply**.

The FTP client configuration is changed and the change is saved to the running configuration.

Configuring the ASA as a Secure Copy Server

You can enable the secure copy server on the ASA. Only clients that are allowed to access the ASA using SSH can establish a secure copy connection.

Restrictions

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 ASA internal files.
- The server does not support banners.
- The server does not support wildcards.
- The ASA license must have the VPN-3DES-AES feature to support SSH Version 2 connections.

To configure the ASA as a secure copy server, perform the following steps:

Detailed Steps

Step 1 From the Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server pane, check the **Enable secure copy server** check box.

Step 2 Click **Apply**.

The changes are saved to the running configuration. The ASA can function as an SCP server for transferring files to and from the device.

Configuring the ASA as a TFTP Client

TFTP is a simple client/server file transfer protocol, which is described in RFC 783 and RFC 1350 Rev. 2. You can configure the ASA as a TFTP *client* so that it can transfer a copy of its running configuration file to a TFTP *server* by choosing either **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 ASAs.

The ASA supports only one TFTP client. The full path to the TFTP client is specified in Configuration > Device Management > Management Access > File Access > TFTP Client. After the TCP client has been configured in this pane, 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 ASA to the TFTP client is done apart from this function.

To configure the ASA 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 the **Enable** check box.

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, enter the IP address of the TFTP server on which configuration files will be saved.

Step 4 In the Path field, enter the path to the TFTP server on which 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 ASA configuration files. For more information, see the "Saving the Running Configuration to a TFTP Server" section on page 101-27.

Adding Mount Points

This section includes the following topics:

- Adding a CIFS Mount Point, page 52-10
- Adding an FTP Mount Point, page 52-11

Adding a CIFS Mount Point

To define a Common Internet File System (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 the **Enable mount point** check box.

This option attaches the CIFS file system on the ASA to the UNIX file tree.

Step 3 In the Mount Point Name field, enter the name of an existing CIFS location.

Step 4 In the Server Name or IP Address field, enter the name or IP address of the server in which the mount point is located.

Step 5 In the Share Name field, enter the name of the folder on the CIFS server.

Step 6 In the NT Domain Name field, enter the name of the NT Domain in which the server resides.

Step 7 In the User Name field, enter the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, enter the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, reenter the password.

Step 10 Click **OK**.

The Add CIFS Mount Point dialog box closes.

Step 11 Click **Apply**.

The mount point is added to the ASA, 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 the 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 ASA to the UNIX file tree.

Step 3 In the Mount Point Name field, enter the name of an existing FTP location.

Step 4 In the Server Name or IP Address field, enter 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 the data connection. The server responds with the number of its listening port for this connection.

Step 6 In the Path to Mount field, enter the directory path name to the FTP file server.

Step 7 In the User Name field, enter the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, enter the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, reenter the password.

Step 10 Click **OK**.

The Add FTP Mount Point dialog box closes.

Step 11 Click **Apply**.

The mount point is added to the ASA, and the change is saved to the running configuration.

Configuring ICMP Access

By default, you can send ICMP packets to any ASA interface using either IPv4 or IPv6. This section tells how to limit ICMP management access to the ASA. You can protect the ASA from attacks by limiting the addresses of hosts and networks that are allowed to have ICMP access to the ASA.

Note For allowing ICMP traffic through the ASA, see Chapter 51, "Configuring Access Rules."

This section includes the following topics:

- Information About ICMP Access, page 52-12
- Licensing Requirements for ICMP Access, page 52-12
- Guidelines and Limitations, page 52-13
- Default Settings, page 52-13
- Configuring ICMP Access, page 52-13

Information About ICMP Access

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.

We recommend that you always grant permission 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 ASA uses a first match to the ICMP traffic followed by an implicit deny all entry. 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 ASA 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.

Licensing Requirements for ICMP Access

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.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

- The ASA does not respond to ICMP echo requests directed to a broadcast address.
- The ASA only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.
- If you cannot ping the ASA interface, make sure that you enable ICMP to the ASA for your IP address using the **icmp** command.

Default Settings

By default, you can send ICMP packets to any ASA interface using either IPv4 or IPv6.

Configuring ICMP Access

To configure ICMP access rules, perform the following steps:

Detailed Steps

Step 1 Choose **Configuration > Device Management > Management Access > ICMP**, and click **Add**.

Step 2 Choose which version of IP traffic to filter by clicking the applicable 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, select 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.

Step 5 From the Interface list, choose the destination ASA interface to which the rule is to be applied.

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,** then go to Step 9.

Step 7 From the Mask drop-down list, choose the network mask.

Step 8 Click **OK**.

The Create ICMP Rule 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 ASA that shows the ASA 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 ASA, and the change is saved to the running configuration.

Configuring Management Access Over a VPN Tunnel

If your VPN tunnel terminates on one interface, but you want to manage the ASA by accessing a different interface, you can identify that interface as a management-access interface. For example, if you enter the ASA from the outside interface, this feature lets you connect to the inside interface using ASDM, SSH, Telnet, or SNMP; or you can ping the inside interface when entering from the outside interface. Management access is available via the following VPN tunnel types: IPsec clients, IPsec site-to-site, and the AnyConnect SSL VPN client.

This section includes the following topics:

- Licensing Requirements for a Management Interface, page 52-14
- Guidelines and Limitations, page 52-2
- Configuring a Management Interface, page 52-15

Licensing Requirements for a Management Interface

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 mode.

Firewall Mode Guidelines

Supported in routed mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

You can define only one management access interface.

Configuring a Management Interface

To configure the management interface, perform the following steps.

Detailed 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 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 procedures listed in Chapter 46, "Configuring AAA Servers and the Local Database."

This section includes the following topics:

- Information About AAA for System Administrators, page 52-16
- Licensing Requirements for AAA for System Administrators, page 52-19
- Prerequisites, page 52-19
- Guidelines and Limitations, page 52-20
- Default Settings, page 52-20
- Configuring Authentication for CLI, ASDM, and enable command Access, page 52-20
- Limiting User CLI and ASDM Access with Management Authorization, page 52-21
- Configuring Command Authorization, page 52-23
- Configuring Management Access Accounting, page 52-29
- Viewing the Currently Logged-In User, page 52-29
- Recovering from a Lockout, page 52-30

Information About AAA for System Administrators

This section describes AAA for system administrators and includes the following topics:

- Information About Management Authentication, page 52-16
- Information About Command Authorization, page 52-17

Information About Management Authentication

This section describes authentication for management access and includes the following topics:

- Comparing CLI Access with and without Authentication, page 52-16
- Comparing ASDM Access with and without Authentication, page 52-16
- Authenticating Sessions from the Switch to the ASA Services Module, page 52-16

Comparing CLI Access with and without Authentication

How you log into the ASA depends on whether or not you enable authentication:

- If you do not enable any authentication for Telnet, you do not enter a username; you enter the login password. For SSH, you enter the username and the login password. You access user EXEC mode.
- If you enable Telnet or SSH authentication according to this section, you enter the username and password as defined on the AAA server or local user database. You access user EXEC mode.

To enter privileged EXEC mode after logging in, enter the **enable** command. How **enable** works depends on whether you enable authentication:

- If you do not configure enable authentication, enter the system enable password when you enter the **enable** 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.
- If you configure enable authentication, the ASA prompts you for your username and password again. This feature is particularly useful when you perform command authorization, in which usernames are important in determining the commands that a user can enter.

For enable authentication using the local database, you can use the **login** command instead of the **enable** command. **login** maintains the username but requires no configuration to turn on authentication.

Comparing ASDM Access with and without Authentication

By default, you can log into ASDM with a blank username and the enable password. Note that 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.

If you configure HTTP authentication, you can no longer use ASDM with a blank username and the enable password.

Authenticating Sessions from the Switch to the ASA Services Module

For sessions from the switch to the ASASM (using the **session** command), you can configure Telnet authentication. For virtual console connections from the switch to the ASASM (using the **service-module session** command), you can configure serial port authentication.

In multiple context mode, you cannot configure any AAA commands in the system configuration. However, if you configure Telnet or serial authentication in the admin context, then authentication also applies to sessions from the switch to the ASASM. The admin context AAA server or local user database is used in this instance.

Information About Command Authorization

This section describes command authorization and includes the following topics:

- Supported Command Authorization Methods, page 52-17
- About Preserving User Credentials, page 52-17
- Security Contexts and Command Authorization, page 52-18

Supported Command Authorization Methods

You can use one of two command authorization methods:

- Local privilege levels—Configure the command privilege levels on the ASA. When a local, RADIUS, or LDAP (if you map LDAP attributes to RADIUS attributes) user authenticates for CLI access, the ASA 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 assigned 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 ASA 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 ASA places you in level *n*. These levels are not used unless you enable local command authorization (see the "Configuring Local Command Authorization" section on page 52-23). (See the command reference for more information about the **enable** command.)

- TACACS+ server privilege levels—On the TACACS+ server, configure the commands that a user or group can use after authenticating for CLI access. Every command that a user enters at the CLI is validated with the TACACS+ server.

About Preserving User Credentials

When a user logs into the ASA, that user is required to provide a username and password for authentication. The ASA 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 for 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.

- The local server is configured to authenticate user access.
- Privilege level 15 command access is configured to require a password.
- The user account is configured for serial-only authorization (no access to console or ASDM).
- The user account is configured for privilege level 15 command access.

The following table shows how credentials are used in this case by the ASA.

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 among contexts.

  When configuring command authorization, you must configure each security context separately. This configuration 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 which 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 that they need.

Note The system execution space does not support AAA commands; therefore, command authorization is not available in the system execution space.

Licensing Requirements for AAA for System Administrators

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites

Depending on the feature, you can use the following:

- AAA server—See the "Configuring AAA Server Groups" section on page 46-11.
- Local Database—See the "Adding a User Account to the Local Database" section on page 46-20.

Prerequisites for Management Authentication

Before the ASA can authenticate a Telnet, SSH, or HTTP user, you must identify the IP addresses that are allowed to communicate with the ASA. For the ASASM, the exception is for access to the system in multiple context mode; a session from the switch to the ASASM is a Telnet session, but Telnet access configuration is not required. For more information, see the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 52-1.

Prerequisites for Local Command Authorization

- Configure **enable** authentication. (See the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20.)

  enable authentication is essential for maintaining 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.
  - 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 VSA CVPN3000-Privilege-Level according to the "Using User Login Credentials" section on page 46-8.

Prerequisites for TACACS+ Command Authorization

- Configure CLI and **enable** authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20).

Prerequisites for Management Accounting

- Configure CLI and **enable** authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 52-20).

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.

Default Settings

Default Command Privilege Levels

By default, the following commands are assigned to privilege level 0. All other commands are assigned to privilege 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.

To view all privilege levels, see the "Viewing Local Command Privilege Levels" section on page 52-24.

Configuring Authentication for CLI, ASDM, and enable command Access

To configure management authentication, perform the following steps:

Detailed Steps

Step 1 To authenticate users who use the **enable** command, choose **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 ASA 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 ASA prompt does not give any indication of which method is being used.

Step 2 To authenticate users who access the CLI or ASDM, choose **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 ASA using HTTPS. HTTP management authentication does not support the SDI protocol for a AAA server group.
- **Serial**—Authenticates users who access the ASA using the console port. For the ASASM, this parameter affects the virtual console accessed from the switch using the **service-module session** command. For multiple mode access, see the "Authenticating Sessions from the Switch to the ASA Services Module" section on page 52-16.
- **SSH**—Authenticates users who access the ASA using SSH.
- **Telnet**—Authenticates users who access the ASA using Telnet. For the ASASM, this parameter also affects the session from the switch using the **session** command. For multiple mode access, see the "Authenticating Sessions from the Switch to the ASA Services Module" section on page 52-16.

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 ASA 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 ASA prompt does not give any indication of 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.

Detailed Steps

Step 1 To enable management authorization, choose **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 52-23 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—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—Request authorization with the "service=shell" entry, 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. By default, the access restriction is Full Access, which allows full access to any services specified by the Authentication tab options. For more information, see the "Adding a User Account to the Local Database" section on page 46-20.

Configuring Command Authorization

If you want to control access to commands, the ASA 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.

You can use one of two command authorization methods:

- Local privilege levels
- TACACS+ server privilege levels

For more information about command authorization, see the "Information About Command Authorization" section on page 52-17.

This section includes the following topics:

- Configuring Local Command Authorization, page 52-23
- Viewing Local Command Privilege Levels, page 52-24
- Configuring Commands on the TACACS+ Server, page 52-25
- Configuring TACACS+ Command Authorization, page 52-28

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 the assigned privilege level or below. The ASA 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 "Using Certificates and User Login Credentials" section on page 46-8.)

To configure local command authorization, perform the following steps:

Detailed Steps

Step 1 To enable command authorization, choose **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 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 or 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 commands that appear, click **Select All** and then **Edit**.

Click **OK** to accept your changes.

Step 4 To support administrative user privilege levels from RADIUS, check the **Perform authorization for exec shell access > Enable** check box.

Without this option, the ASA 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 52-21 for more information.

Step 5 Click **Apply**.

The authorization settings are assigned, and the changes are saved to the running configuration.

Viewing Local Command Privilege Levels

The following commands when entered in the Tools > Command Line Interface tool, let you view privilege levels for commands.

Examples

For the **show running-config all privilege all** command, the ASA displays the current assignment of each CLI command to a privilege level. The following is sample output from this command:

Enter the following command in the Tools > Command Line Interface tool:

```
show running-config all privilege all
privilege show level 15 command aaa
privilege clear level 15 command aaa
privilege configure level 15 command aaa
privilege show level 15 command aaa-server
privilege clear level 15 command aaa-server
privilege configure level 15 command aaa-server
privilege show level 15 command access-group
privilege clear level 15 command access-group
privilege configure level 15 command access-group
privilege show level 15 command access-list
privilege clear level 15 command access-list
```

```
privilege configure level 15 command access-list
privilege show level 15 command activation-key
privilege configure level 15 command activation-key
....
```

The following example displays the command assignments for privilege level 10:

```
show running-config privilege level 10
privilege show level 10 command aaa
```

The following example displays the command assignments for the **access-list** command:

```
show running-config privilege command access-list
privilege show level 15 command access-list
privilege clear level 15 command access-list
privilege configure level 15 command access-list
```

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 ASA 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 ASA 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 field, and type **permit aaa-server** in the arguments field.

- You can permit all arguments of a command that you do not explicitly deny by checking 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 a question mark, which shows CLI usage (see Figure 52-1).

Figure 52-1 Permitting All Related Commands

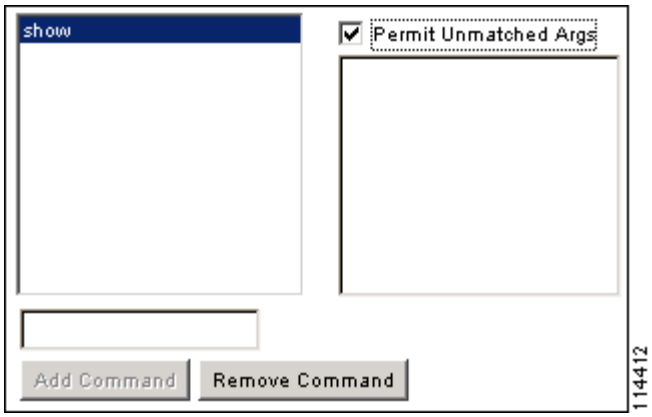

- For commands that are a single word, you *must* permit unmatched arguments, even if there are no arguments for the command, for example **enable** or **help** (see Figure 52-2).

Figure 52-2 Permitting Single Word Commands

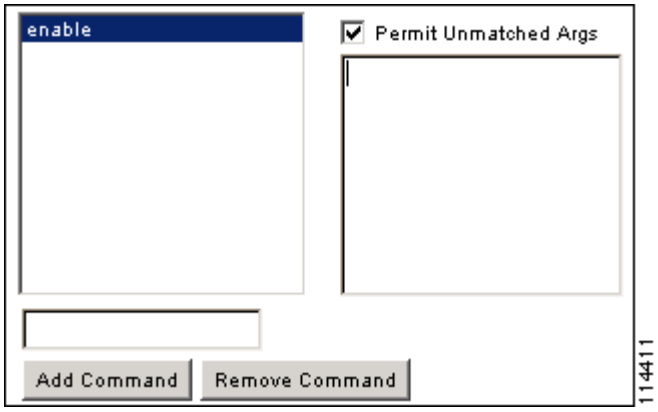

- To disallow some arguments, enter the arguments preceded by **deny**.

  For example, to allow **enable**, but not **enable password**, enter **enable** in the commands field, and **deny password** in the arguments field. Be sure to check the **Permit Unmatched Args** check box so that **enable** alone is still allowed (see Figure 52-3).

Figure 52-3 Disallowing Arguments

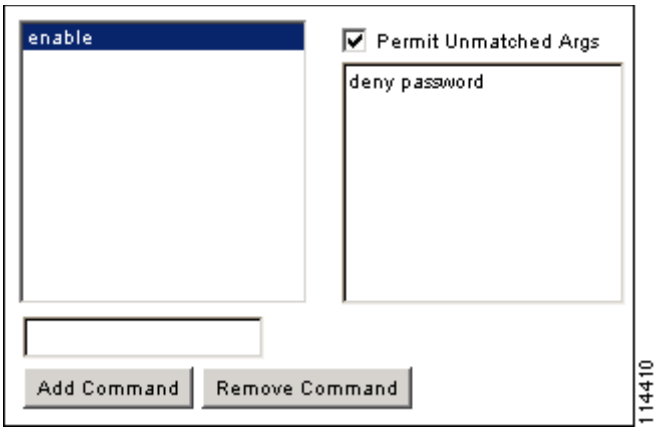

- When you abbreviate a command at the command line, the ASA expands the prefix and main command to the full text, but it sends additional arguments to the TACACS+ server as you enter them.

  For example, if you enter **sh log**, then the ASA sends the entire command to the TACACS+ server, **show logging**. However, if you enter **sh log mess**, then the ASA 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 52-4).

Figure 52-4 Specifying Abbreviations

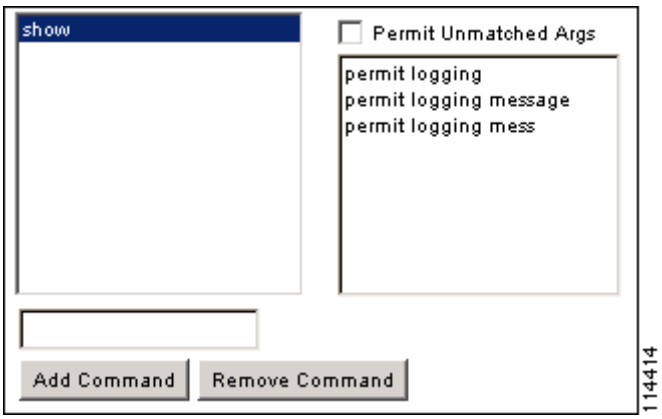

- We recommend that you allow the following basic commands for all users:
  - **show checksum**
  - **show curpriv**
  - **enable**
  - **help**
  - **show history**
  - **login**
  - **logout**
  - **pager**

- **show pager**
- **clear pager**
- **quit**
- **show version**

Configuring TACACS+ Command Authorization

If you enable TACACS+ command authorization, and a user enters a command at the CLI, the ASA sends the command and username to the TACACS+ server to determine if the command is authorized.

Before you enable TACACS+ command authorization, be sure that you are logged into the ASA as a user that is defined on the TACACS+ server, and that you have the necessary command authorization to continue configuring the ASA. For example, you should log in as an admin user with all commands authorized. Otherwise, you could become unintentionally locked out.

Do not save your configuration until you are sure that it works the way you want. If you get locked out because of a mistake, you can usually recover access by restarting the ASA. If you still get locked out, see the "Recovering from a Lockout" section on page 52-30.

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 ASA. 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 procedures listed in the "Configuring Command Authorization" section on page 52-23.

To configure TACACS+ command authorization, perform the following steps:

Detailed Steps

Step 1 To perform command authorization using a TACACS+ server, choose **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 ASA to use the local database as a fallback method if the AAA server is unavailable. To do so, check 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 ASA prompt does not give any indication which method is being used. Be sure to configure users in the local database (see the "Adding a User Account to the Local Database" section on page 46-20) and command privilege levels (see the "Configuring Local Command Authorization" section on page 52-23).

Step 4 Click **Apply**.

The command authorization settings are assigned, and the changes are saved to the running configuration.

Configuring Management Access Accounting

You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI. You can configure accounting when users log in, when they enter the **enable** command, or when they issue commands.

For command accounting, you can only use TACACS+ servers.

To configure management access and enable command accounting, perform the following steps:

Detailed Steps

Step 1 To enable accounting of users when they enter the **enable** command, perform the following steps:

a. Choose **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 2 To enable accounting of users when they access the ASA using Telnet, SSH, or the serial console, perform the following steps:

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 3 To configure command accounting, perform the following steps:

a. Under the Require command accounting area, check the **Enable** check box.

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, you can limit which commands the ASA accounts for by specifying a minimum privilege level in the Privilege level drop-down list. The ASA does not account for commands that are below the minimum privilege level.

Step 4 Click **Apply**.

The accounting settings are assigned, and the changes are saved to the running configuration.

Viewing the Currently Logged-In User

To view the current logged-in user, enter the following command in the Tools > Command Line Interface tool:

```
show curpriv
```

The following is sample output from the **show curpriv** command:

```
show curpriv
Username: admin
Current privilege level: 15
```

```
Current Mode/s: P_PRIV
```

Table 52-1 describes the **show curpriv** command output.

Table 52-1 show curpriv Command Output Description

Field	Description
Username	Username. If you are logged in as the default user, the name is enable_1 (user EXEC) or enable_15 (privileged EXEC).
Current privilege level	Levels range from 0 to 15. Unless you configure local command authorization and assign commands to intermediate privilege levels, levels 0 and 15 are the only levels that are used.
Current Mode/s	The available access modes are the following: • P_UNPR—User EXEC mode (levels 0 and 1) • P_PRIV—Privileged EXEC mode (levels 2 to 15) • P_CONF—Configuration mode

Recovering from a Lockout

In some circumstances, when you turn on command authorization or CLI authentication, you can be locked out of the ASA CLI. You can usually recover access by restarting the ASA. However, if you already saved your configuration, you might be locked out. Table 52-2 lists the common lockout conditions and how you might recover from them.

Table 52-2 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 ASA 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 ASA, session into the ASA 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 52-2 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 ASA immediately, then log into the maintenance partition and reset the passwords and **aaa** commands.	Session into the ASA 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 ASA from the switch. From the system execution space, you can change to the context and change the user level.

Monitoring Device Access

To monitor device access, see the following panes:

Path	Purpose
Monitoring > Properties > Device Access > ASDM/HTTPS/Telnet/SSH Sessions	The top pane lists the connection types, session IDs, and IP addresses for users connected through ASDM, HTTPS, and Telnet sessions. To disconnect a specific session, click **Disconnect**. The bottom pane lists the clients, usernames, connection states, software versions, incoming encryption types, outgoing encryption types, incoming HMACs, and outgoing HMACs for users connected through SSH sessions. To disconnect a specific session, click **Disconnect**.
`Monitoring > Properties > Device Access > Authenticated Users`	Lists the usernames, IP addresses, dynamic ACLs, inactivity timeouts (if any), and absolute timeouts for users who were authenticated by AAA servers.
`Monitoring > Properties > Device Access > AAA Local Locked Out Users`	Lists the usernames of locked-out AAA local users, the number of failed attempts to authenticate, and the times that users were locked out. To clear a specific user who has been locked out, click **Clear Selected Lockout**. To clear all users who have been locked out, click **Clear All Lockouts**.

Feature History for Management Access

Table 52-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 52-3 Feature History for Management Access

Feature Name	Platform Releases	Feature Information
Management Access	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH Configuration > Device Management > Management Access > Command Line (CLI) > Banner Configuration > Device Management > Management Access > CLI Prompt Configuration > Device Management > Management Access > ICMP Configuration > Device Management > Management Access > File Access > FTP Client Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server Configuration > Device Management > Management Access > File Access > Mount-Points Configuration > Device Management > Users/AAA > AAA Access > Authentication Configuration > Device Management > Users/AAA > AAA Access > Authorization Configuration > Device Management > Users/AAA > AAA Access > > Accounting.
Increased SSH security; the SSH default username is no longer supported.	8.4(2)	Starting in 8.4(2), you can no longer connect to the ASA using SSH with the pix or asa username and the login password. To use SSH, you must configure AAA authentication using the **aaa authentication ssh console LOCAL** command (CLI) or Configuration > Device Management > Users/AAA > AAA Access > Authentication (ASDM); then define a local user by entering the **username** command (CLI) or choosing Configuration > Device Management > Users/AAA > User Accounts (ASDM). If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.
For the ASASM in multiple context mode, support for Telnet and virtual console authentication from the switch.	8.5(1)	Although connecting to the ASASM from the switch in multiple context mode connects to the system execution space, you can configure authentication in the admin context to govern those connections.

Configuring AAA Rules 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 52-15.

This chapter includes the following sections:

- AAA Performance, page 53-1
- Licensing Requirements for AAA Rules, page 53-1
- Guidelines and Limitations, page 53-2
- Configuring Authentication for Network Access, page 53-2
- Configuring Authorization for Network Access, page 53-11
- Configuring Accounting for Network Access, page 53-17
- Using MAC Addresses to Exempt Traffic from Authentication and Authorization, page 53-18
- Feature History for AAA Rules, page 53-19

AAA Performance

The ASA 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 ASA cut-through proxy challenges a user initially at the application layer and then authenticates with standard AAA servers or the local database. After the ASA 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.

Licensing Requirements for AAA Rules

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.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

In clustering, this feature is only supported on the master unit.

Configuring Authentication for Network Access

This section includes the following topics:

- Information About Authentication, page 53-2
- Configuring Network Access Authentication, page 53-6
- Enabling the Redirection Method of Authentication for HTTP and HTTPS, page 53-7
- Enabling Secure Authentication of Web Clients, page 53-8
- Authenticating Directly with the ASA, page 53-9
- Configuring the Authentication Proxy Limit, page 53-11

Information About Authentication

The ASA lets you configure network access authentication using AAA servers. This section includes the following topics:

- One-Time Authentication, page 53-3
- Applications Required to Receive an Authentication Challenge, page 53-3
- ASA Authentication Prompts, page 53-3
- AAA Prompts and Identity Firewall, page 53-4
- AAA Rules as a Backup Authentication Method, page 53-5
- Static PAT and HTTP, page 53-5

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 ASA 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 ASA 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 ASA allows other traffic requiring authentication.

The authentication ports that the ASA supports for AAA are fixed as follows:

- Port 21 for FTP
- Port 23 for Telnet
- Port 80 for HTTP
- Port 443 for HTTPS

ASA Authentication Prompts

For Telnet and FTP, the ASA generates an authentication prompt.

For HTTP, the ASA uses basic HTTP authentication by default, and provides an authentication prompt. You can optionally configure the ASA to redirect users to an internal web page where they can enter their username and password (configured in 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 53-7).

For HTTPS, the ASA generates a custom login screen. You can optionally configure the ASA to redirect users to an internal web page where they can enter their username and password (configured in 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 53-7).

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 ASA.

You might want to continue to use basic HTTP authentication for the following reasons:

- You do not want the ASA to open listening ports.
- You use NAT on a router and you do not want to create a translation rule for the web page served by the ASA.
- Basic HTTP authentication might work better with your network.

For example non-browser applications, as when a URL is embedded in e-mail, might be more compatible with basic authentication.

After you authenticate correctly, the ASA 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 ASA 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 53-8 for information to secure your credentials.

For FTP, a user has the option of entering the ASA username followed by an at sign (@) and then the FTP username (name1@name2). For the password, the user enters the ASA password followed by an at sign (@) and then the FTP password (password1@password2). For example, enter the following text:

```
name> name1@name2
password> password1@password2
```

This feature is useful when you have cascaded firewalls that require multiple logins. You can separate several names and passwords by multiple at signs (@).

AAA Prompts and Identity Firewall

In an enterprise, some users log into the network by using other authentication mechanisms, such as authenticating with a web portal (cut-through proxy). For example, users with a Mac and Linux client might log into a web portal (cut-through proxy). Therefore, you must configure the identity firewall to allow these types of authentication in connection with identity-based access policies.

Figure 53-1 shows a deployment to support a cut-through proxy authentication captive portal. Active Directory servers and the AD Agent are installed on the main site LAN. However, the identity firewall is configured to support authentication of clients that are not part of the Active Directory domain.

Figure 53-1 Deployment Supporting Cut-through Proxy Authentication

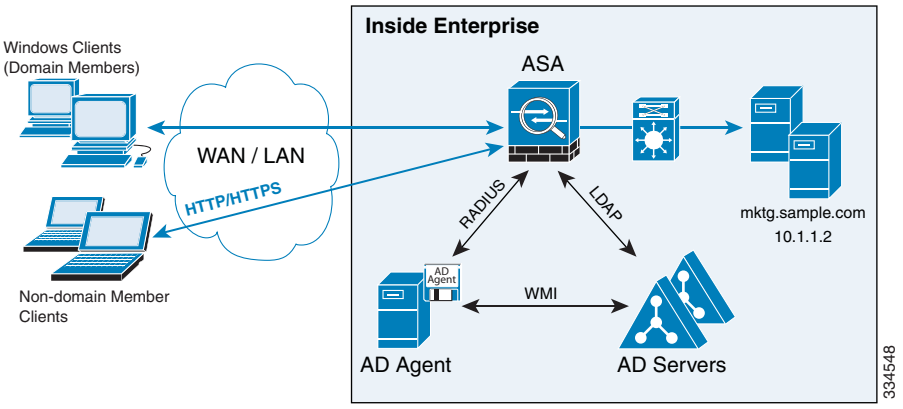

The ASA designates users logging in through a web portal (cut-through proxy) as belonging to the Active Directory domain with which they authenticated.

The ASA reports users logging in through a web portal (cut-through proxy) to the AD Agent, which distributes the user information to all registered ASA devices. In this case, the identity firewall can associate the users with their Active Directory domain. Specifically, the user identity-IP address mappings of authenticated users are forwarded to all ASA contexts that contain the input interface where packets are received and authenticated.

Users can log in by using HTTP/HTTPS, FTP, Telnet, or SSH. When users log in with these authentication methods, the following guidelines apply:

- For HTTP/HTTPS traffic, an authentication window appears for unauthenticated users.

- For Telnet and FTP traffic, users must log in through the cut-through proxy server and again to the Telnet and FTP servers.
- A user can specify an Active Directory domain while providing login credentials (in the format, domain\username). The ASA automatically selects the associated AAA server group for the specified domain.
- If a user specifies an Active Directory domain while providing login credentials (in the format, domain\username), the ASA parses the domain and uses it to select an authentication server from the AAA servers that have been configured for the identity firewall. Only the username is passed to the AAA server.
- If the backslash (\) delimiter is not found in the login credentials, the ASA does not parse the domain and authentication is conducted with the AAA server that corresponds to the default domain configured for the identity firewall.
- If a default domain or a server group is not configured for that default domain, the ASA rejects the authentication.
- If the domain is not specified, the ASA selects the AAA server group for the default domain that is configured for the identity firewall.

AAA Rules as a Backup Authentication Method

An authentication rule (also known as "cut-through proxy") controls network access based on the user. Because this function is very similar to an access rule plus an identity firewall, AAA rules can now be used as a backup method of authentication if a user's AD login expires or a valid user has not yet logged into AD. For example, for any user without a valid login, you can trigger a AAA rule. To ensure that the AAA rule is only triggered for users that do not have valid logins, you can specify special usernames in the extended ACL that are used for the access rule and for the AAA rule: None (users without a valid login) and Any (users with a valid login). In the access rule, configure your policy as usual for users and groups, but then include a rule that permits all None users before deny any any; you must permit these users so they can later trigger a AAA rule. Then, configure a AAA rule that does not match Any users (these users are not subject to the AAA rule, and were handled already by the access rule), but matches all None users only to trigger AAA authentication for these users. After the user has successfully logged in via cut-through proxy, the traffic will flow normally again.

Static PAT and HTTP

For HTTP authentication, the ASA checks real ports when static PAT is configured. If it detects traffic destined for real port 80, regardless of the mapped port, the ASA intercepts the HTTP connection and enforces authentication.

For example, assume that outside TCP port 889 is translated to port 80 and that any relevant access lists permit the traffic:

```
object network obj-192.168.123.10-01
    host 192.168.123.10
    nat (inside,outside) static 10.48.66.155 service tcp 80 889
```

Then when users try to access 10.48.66.155 on port 889, the ASA intercepts the traffic and enforces HTTP authentication. Users see the HTTP authentication page in their web browsers before the ASA allows HTTP connection to complete.

If the local port is different than port 80, as in the following example:

```
object network obj-192.168.123.10-02
    host 192.168.123.10
```

```
nat (inside,outside) static 10.48.66.155 service tcp 111 889
```

Then users do not see the authentication page. Instead, the ASA sends an error message to the web browser indicating that the user must be authenticated before using the requested service.

Configuring Network Access Authentication

To configure network access authentication, perform the following steps:

Step 1 In the Configuration > Firewall > AAA Rules pane, choose **Add > Add Authentication Rule**.

The Add Authentication Rule dialog box appears.

Step 2 In the Interface drop-down list, choose the interface for applying the rule.

Tip In the Action field, click one of the following, depending on the implementation:

- **Authenticate**
- **Do not Authenticate**

Step 3 In 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 "Testing Server Authentication and Authorization" section on page 46-27 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 to the Local Database" section on page 46-20 for more information.

Step 4 In the Source field, add the source IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 5 In the Destination field, enter the destination IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 6 In the Service field, enter an IP service name or number for the destination service, or click the ellipsis (...) to choose a service.

Step 7 (Optional) In the Description field, enter a description.

Step 8 (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, clear the **Enable Rule** check box.

  You may not want to remove a rule, but instead turn it off.
- To set a time range for the rule, In the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see the "Configuring Time Ranges" section on page 25-15.

Step 9 Click **OK**.

The Add Authentication Rule dialog box closes and the rule appears in the AAA Rules table.

Step 10 Click **Apply**.

The changes are saved to the running configuration.

For more information about authentication, see the "Information About Authentication" section on page 53-2.

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 ASA 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 "ASA Authentication Prompts" section on page 53-3 for more information about the redirection method compared to the basic method.

To enable a AAA listener, perform the following steps:

Step 1 In 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 In the Interface drop-down list, choose the interface on which you want to enable the listener.

Step 5 In the Port drop-down list, choose the port or enter a number.

This is the port that the ASA 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 ASA. Without this option, only traffic directed to the ASA 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** again to close the AAA Rules Advanced Options dialog box.

Step 8 Click **Apply**.

The changes are saved to the running configuration.

Enabling Secure Authentication of Web Clients

If you use HTTP authentication, by default the username and password are sent from the client to the ASA in clear text; in addition, the username and password are sent to the destination web server as well.

The ASA provides the following methods for securing HTTP authentication:

- Enable the redirection method of authentication for HTTP—See the "Enabling the Redirection Method of Authentication for HTTP and HTTPS" section on page 53-7. This method prevents the authentication credentials from continuing to the destination server. See the "ASA Authentication Prompts" section on page 53-3 for more information about the redirection method compared to the basic method.
- Enable virtual HTTP— Virtual HTTP lets you authenticate separately with the ASA and with the HTTP server. Even if the HTTP server does not need a second authentication, this command 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 53-9 for more information.
- Enable the exchange of usernames and passwords between a web client and the ASA with HTTPS—To enable the exchange of usernames and passwords between a web client and the ASA with HTTPS, perform the following steps:

  a. In 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** again to close the AAA Rules Advanced Options dialog box.

  d. Click **Apply**.

  This is the only method that protects credentials between the client and the ASA, as well as between the ASA 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 ASA redirects the HTTP user to an HTTPS prompt. After you authenticate correctly, the ASA 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 one second (see the Configuration > Firewall > Advanced > Global Timeouts pane). However, this workaround opens a 1-second window of opportunity that might allow unauthenticated 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 ruleto block traffic from the HTTP client to the 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 ASA

If you do not want to allow HTTP, HTTPS, Telnet, or FTP through the ASA but want to authenticate other types of traffic, you can authenticate with the ASA directly using HTTP, HTTPS, or Telnet.

This section includes the following topics:

- Authenticating HTTP(S) Connections with a Virtual Server, page 53-9
- Authenticating Telnet Connections with a Virtual Server, page 53-10

Authenticating HTTP(S) Connections with a Virtual Server

If you enabled the redirection method of HTTP and HTTPS authentication in the "Configuring Network Access Authentication" section on page 53-6, then you have also automatically enabled direct authentication.

When you use HTTP authentication on the ASA (see the"Configuring Network Access Authentication" section on page 53-6), the ASA uses basic HTTP authentication by default.

You can change the authentication method so that the ASA redirects HTTP connections to web pages generated by the ASA itself using the "Enabling the Redirection Method of Authentication for HTTP and HTTPS" section on page 53-7.

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 ASA, then virtual HTTP lets you authenticate separately with the ASA (via a AAA server) and with the HTTP server. Without virtual HTTP, the same username and password that you used to authenticate with the ASA is sent to the HTTP server; you are not prompted separately for the HTTP server username and password. Assuming the username and password are 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 ASA. The ASA prompts for the AAA server username and password. After the AAA server authenticates the user, the ASA 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 71-9.

You can authenticate directly with the ASA 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 ASA virtual server separately from the HTTP server, perform the following steps:

Step 1 In 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 ASA. 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.

Authenticating Telnet Connections with a Virtual Server

Although you can configure network access authentication for any protocol or service(see the "Configuring Network Access Authentication" section on page 53-6), 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 traffic through the ASA, but want to authenticate other types of traffic, you can configure virtual Telnet; the user Telnets to a given IP address configured on the ASA, and the ASA issues a Telnet prompt.

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. After the user is authenticated, the message "Authentication Successful" appears. 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. In addition, 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 log out from the ASA, 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 In the Configuration > Firewall > Advanced > Virtual Access > Virtual Telnet Server area, check the **Enable** check box.

Step 2 In the Virtual Telnet Server field, enter the IP address of the virtual Telnet server.

Make sure that this address is an unused address that is routed to the ASA. 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.

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 Choose **Configuration > Firewall > AAA Rules**, then click **Advanced**.

The AAA Rules Advanced Options dialog box appears.

Step 2 In the Proxy Limit area, check the **Enable Proxy Limit** check box.

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**, then click **Apply**.

The changes are saved to the running configuration.

Configuring Authorization for Network Access

After a user authenticates for a given connection, the ASA can use authorization to further control traffic from the user.

This section includes the following topics:

- Configuring TACACS+ Authorization, page 53-11
- Configuring RADIUS Authorization, page 53-13

Configuring TACACS+ Authorization

You can configure the ASA to perform network access authorization with TACACS+. Authentication and authorization statements 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 ASA.

   Because a user at a given IP address only needs to authenticate one time for all rules and types, if the authentication session has not expired, authorization can occur even if the traffic is not matched by an authentication rule.
2. After a user authenticates, the ASA checks the authorization rules for matching traffic.
3. If the traffic matches the authorization rule, the ASA sends the username to the TACACS+ server.
4. The TACACS+ server responds to the ASA with a permit or a deny for that traffic, based on the user profile.
5. The ASA 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 53-6. If you have already enabled authentication, continue to the next step.

Step 2 In the Configuration > Firewall > AAA Rules pane, choose **Add > Add Authorization Rule**.

The Add Authorization Rule dialog box appears.

Step 3 In 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 In 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 "Testing Server Authentication and Authorization" section on page 46-27 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 the ellipsis (...) to choose a service.

Step 9 (Optional) In the Description field, enter 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 content into the Source Service field.
- To make the rule inactive, clear the **Enable Rule** check box.

  You may not want to remove a rule, but instead turn it off.
- To set a time range for the rule, in the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see the "Configuring Time Ranges" section on page 25-15.

Step 11 Click **OK**.

The Add Authorization Rule 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 Network Access Authentication" section on page 53-6.

When you configure the ASA 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 ASA. It does provide information about how the ASA handles access list information received from RADIUS servers.

You can configure a RADIUS server to download an access list to the ASA 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 the per-user-override featureon 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 53-13
- Configuring a RADIUS Server to Download Per-User Access Control List Names, page 53-17

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 53-13
- Configuring Cisco Secure ACS for Downloadable Access Lists, page 53-15
- Configuring Any RADIUS Server for Downloadable Access Lists, page 53-16
- Converting Wildcard Netmask Expressions in Downloadable Access Lists, page 53-17

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 ASA.
- 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 ASAs.

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 ASA receives downloadable access lists from Cisco Secure ACS using the following process:

1. The ASA 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 includes the internal name of the applicable downloadable access list. The Cisco IOS cisco-av-pair RADIUS VSA (vendor 9, attribute 1) includes 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 ASA examines the name of the downloadable access list and determines if it has previously received the named downloadable access list.
   - If the ASA has previously received the named downloadable access list, communication with Cisco Secure ACS is complete and the ASA applies the access list to the user session. Because the name of the downloadable access list includes the date and time that it was last modified, matching the name sent by Cisco Secure ACS to the name of an access list previously downloaded means that the ASA has the most recent version of the downloadable access list.
   - If the ASA 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 ASA 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 ASA signs the request with the Message-Authenticator attribute (IETF RADIUS attribute 80).
4. After receipt of a RADIUS authentication request that has a username attribute that includes 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 that includes the access list. The largest access list that can fit in a single access-accept message is slightly less than 4 KB, because part 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 include an ACE and are numbered serially:

   ```
   ip:inacl#1=ACE-1
   ip:inacl#2=ACE-2
   .
   .
   .
   ip:inacl#n=ACE-n

   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 includes a portion of the access list, formatted as described previously, and a State attribute (IETF RADIUS attribute 24), which includes 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 ASA stores the portion of the access list received and responds with another access-request message that includes 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 process 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 ASA commands that are similar to the extended **access-list** command (see command reference), 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 ASA, 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 ASA 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 ASA 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 (see command reference), 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 ASA. 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 ASA, 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 ASA 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 ASA 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 ASA, you may need the ASA to convert wildcard netmask expressions to standard netmask expressions. This is because Cisco VPN 3000 series concentrators support wildcard netmask expressions, but the ASA only supports standard netmask expressions. Configuring the ASA to convert wildcard netmask expressions helps minimize the effects of these differences on 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 ASA 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 in the Configuration > Device Management > Users/AAA > AAA Server Groups > AAA Server Groups area. For more information, see the "Configuring AAA Server Groups" section on page 46-11.

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 ASA 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 values for filter-id attributes are specified in boxes in the HTML interface, omitting **filter-id=** and entering only *acl_name*.

For information about making the filter-id attribute value unique per user, see the documentation for your RADIUS server.

To create an access list on the ASA, see Chapter 26, "Using the ACL Manager."

Configuring Accounting for Network Access

The ASA can send accounting information to a RADIUS or TACACS+ server about any TCP or UDP traffic that passes through the ASA. 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 session start and stop times, username, the number of bytes that pass through the ASA 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 ASA to provide accounting data per user, you must enable authentication. For more information, see the "Configuring Network Access Authentication" section on page 53-6. If you want the ASA to provide accounting data per IP address, enabling authentication is not necessary and you can continue to the next step.

Step 2 In the Configuration > Firewall > AAA Rules pane, choose **Add > Add Accounting Rule**.

The Add Accounting Rule dialog box appears.

Step 3 In 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 In 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 "Testing Server Authentication and Authorization" section on page 46-27 for more information.

Step 6 In the Source field, enter 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 the ellipsis (...) to choose a service.

Step 9 (Optional) In the Description field, enter 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 content to the Source Service field.
- To make the rule inactive, clear the **Enable Rule** check box.

  You may not want to remove a rule, but instead turn it off.
- To set a time range for the rule, In the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see the "Configuring Time Ranges" section on page 25-15.

Step 11 Click **OK**.

The Add Accounting Rule 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 ASA can exempt from authentication and authorization any traffic from specific MAC addresses.

For example, if the ASA 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, instead of 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 1 In the Configuration > Firewall > AAA Rules pane, choose **Add > Add MAC Exempt Rule**.

The Add MAC Exempt Rule dialog box appears.

Step 2 In the Action drop-down list, click one of the following options, 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 3 In the MAC Address field, specify the source MAC address in 12-digit hexadecimal form; that is, nnnn.nnnn.nnnn.

Step 4 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 5 Click **OK**.

The Add MAC Exempt Rule dialog box closes and the rule appears in the AAA Rules table.

Step 6 Click **Apply**.

The changes are saved to the running configuration.

Feature History for AAA Rules

Table 53-1 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 53-1 Feature History for AAA Rules

Feature Name	Platform Releases	Feature Information
AAA Rules	7.0(1)	AAA Rules describe how to enable AAA for network access. We introduced the following screens: Configuration > Firewall > AAA Rules Configuration > Firewall > Advanced > Virtual Access.
Authentication using Cut-Through Proxy	9.0(1)	You can authenticate using AAA rules in conjunction with the Identity Firewall feature.

Configuring Web Cache Services Using WCCP

This chapter describes how to configure web caching services using WCCP, and includes the following sections:

- Information About WCCP, page 54-1
- Guidelines and Limitations, page 54-1
- Licensing Requirements for WCCP, page 54-3
- Adding or Editing WCCP Service Groups, page 54-3
- Configuring Packet Redirection, page 54-4
- WCCP Monitoring, page 54-4
- Feature History for WCCP, page 54-5

Information About WCCP

The purpose of web caching is to reduce latency and network traffic. Previously-accessed web pages are stored in a cache buffer, so if users need the page again, they can retrieve it from the cache instead of the web server.

WCCP specifies interactions between the ASA and external web caches. The feature transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times. The ASA only supports WCCP Version 2.

Using an ASA as an intermediary eliminates the need for a separate router to do the WCCP redirection, because the ASA redirects requests to cache engines. When the ASA determines that a packet needs redirection, it skips TCP state tracking, TCP sequence number randomization, and NAT on these traffic flows.

Guidelines and Limitations

The following WCCPv2 features are supported for the ASA:

- Redirection of multiple TCP and UDP port-destined traffic.
- Authentication for cache engines in a service group.
- Multiple cache engines in a service group.
- GRE encapsulation.

The following WCCPv2 features are not supported for the ASA:

- Multiple routers in a service group.
- Multicast WCCP.
- The Layer 2 redirect method.
- WCCP source address spoofing.
- WAAS devices.

WCCP Interaction With Other Features

In the ASA implementation of WCCP, the protocol interacts with other configurable features according to the following:

- Cut-through proxy will not work in combination with WCCP.
- An ingress access list entry always takes higher priority over WCCP. For example, if an access list does not permit a client to communicate with a server, then traffic is not redirected to a cache engine. Both ingress interface access lists and egress interface access lists are applied.
- TCP intercept, authorization, URL filtering, inspect engines, and IPS features are not applied to a redirected flow of traffic.
- When a cache engine cannot service a request and a packet is returned, or when a cache miss happens on a cache engine and it requests data from a web server, then the contents of the traffic flow is subject to all the other configured features of the ASA.
- If you have two WCCP services and they use two different redirection ACLs that overlap and match the same packets (with a deny or a permit action), the packets behave according to the first service-group found and installed rules. The packets are not passed thorugh all service-groups.

Failover Guidelines

Supports Active/Active and Active/Standby failover. WCCP redirect tables are not replicated to standby units. After a failover, packets are not redirected until the tables are rebuilt. Sessions redirected before failover are probably reset by the web server.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

Context Mode Guidelines

Supported in single mode and multiple context mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

The ASA selects the highest IP address configured on any interface as the WCCP router ID. This address is used to establish a GRE tunnel with the cache engine.

WCCP does not support access lists that include a user, user group, or a fully qualified domain name object.

Licensing Requirements for WCCP

Table 54-1 shows the licensing requirements for WCCP.

Table 54-1 Licensing Requirements

Model	License Requirement
All models	Base License.

Configuring WCCP Service Groups

To allocate space and enable support of the specified WCCP service group, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Device Management > Advanced > WCCP > Service Groups**.

Step 2 To add a new service group, click **Add** to display the Add Service Group dialog box.

Step 3 To modify an existing service group, click **Edit** to display the Edit Service Group dialog box.

Step 4 To remove a selected service group, click **Delete**.

Step 5 To continue, see "Adding or Editing WCCP Service Groups" section on page 54-3.

Step 6 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Adding or Editing WCCP Service Groups

To add a new service group or change the service group parameters for a configured service group, perform the following steps:

Step 1 Click either the Web Cache Service or the Dynamic Service Number radio button. The maximum number of services, including those specified with a dynamic service identifier is 256.

Step 2 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.

Step 3 In the Options area, perform the following steps:

- **a.** Choose the predefined access list that controls traffic redirected to this service group.
- **b.** Choose the predefined access list that determines which web caches are allowed to participate in the service group. Only extended ACLs are allowed.
- **c.** Enter a password up to seven characters long, which is used for MD5 authentication for messages received from the service group.
- **d.** Confirm the password.
- **e.** Click **Manage** to display the ACL Manager window, where you can create or change an ACL.

Step 4 Click **OK** to close the Add or Edit Service Group dialog box.

Step 5 To continue, see the "Configuring Packet Redirection" section on page 54-4.

Configuring Packet Redirection

To configure packet redirection on the ingress of an interface using WCCP, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Device Management > Advanced > WCCP > Redirection**.

Step 2 To add a new WCCP packet redirection, click **Add** to display the Add WCCP Redirection dialog box.

Step 3 To modify an existing WCCP packet redirection, click **Edit** to display the Edit WCCP Redirection dialog box.

Step 4 To remove a selected WCCP packet redirection, click **Delete**.

Step 5 To continue, see the "Adding or Editing Packet Redirection" section on page 54-4.

Adding or Editing Packet Redirection

To add or change packet redirection on the ingress of an interface using WCCP, perform the following steps:

Step 1 Choose the interface on which to enable WCCP redirection from the drop-down list.

Step 2 Choose the service group from the drop-down list.

Step 3 Click **OK** to close the Edit WCCP Redirection dialog box.

Step 4 (Optional) If you need to create a new service group, click **New** to display the Add Service Group dialog box.

Step 5 (Optional) To continue, see the "Adding or Editing WCCP Service Groups" section on page 54-3.

WCCP Monitoring

To monitor WCCP, perform the following steps:

Path	Purpose
Tools > Command Line Interface Type **show running-config wccp**, then click **Send**.	Shows the current WCCP configuration.
Tools > Command Line Interface Type **show running-config wccp interface**, then click **Send**.	Shows the current WCCP interfaces status.

Path	Purpose
Monitoring > Properties > WCCP > WCCP Service Groups	Shows configured WCCP service groups.
Monitoring > Properties > WCCP > WCCP Redirection	Shows configured WCCP interface statistics.

Feature History for WCCP

Table 54-2 lists the release history for this feature. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 54-2 Feature History for WCCP

Feature Name	Releases	Feature Information
WCCP	7.2(1)	WCCP specifies interactions between the ASA and external web caches. We introduced the following screens: **Configuration > Device Management > Advanced > WCCP > Service Groups** **Configuration > Device Management > Advanced > WCCP > Redirection**

Configuring Public Servers

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

- Information About Public Servers, page 55-1
- Licensing Requirements for Public Servers, page 55-1
- Guidelines and Limitations, page 55-1
- Adding a Public Server that Enables Static NAT, page 55-2
- Adding a Public Server that Enables Static NAT with PAT, page 55-2
- Editing Settings for a Public Server, page 55-3
- Feature History for Public Servers, page 55-4

Information About Public Servers

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

Licensing Requirements for Public Servers

Model	License Requirement
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.

Adding a Public Server that Enables Static NAT

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

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

The Add Public Server dialog box appears.

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

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

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

Optionally, from the Browse Service dialog box you can click **Add** to create a new service or service group. Multiple services from various ports can be opened to the outside. For more information about service objects and service groups, see the "Configuring Service Objects and Service Groups" section on page 25-4.

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

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

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

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

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

For information about static NAT, see the "Information About Static NAT" section on page 39-3.

Adding a Public Server that Enables Static NAT with PAT

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

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

The Add Public Server dialog box appears.

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

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

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

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

Optionally, from the Browse Service dialog box, click **Add** to create a new service or service group. Multiple services from various ports can be opened to the outside. For more information about service objects and service groups, see the "Configuring Service Objects and Service Groups" section on page 25-4.

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

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

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

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

Step 10 Click **OK**.

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

For information about static NAT with port address translation, see the "Information About Static NAT with Port Translation" section on page 39-4.

Editing Settings for a Public Server

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

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

The Edit Public Server dialog box appears.

Step 2 Make any necessary changes to the following settings:

- Private Interface—The interface to which the real server is connected.
- Private IP Address—The real IP address of the server.
- Private Service—The actual service that is running on the real server.
- Public Interface—The interface through which outside users can access the real server.
- Public Address—The IP address that is seen by outside users.
- Public Service—The service that is running on the translated address.

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

Feature History for Public Servers

Table 55-1 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 55-1 Feature History for Public Servers

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

PART 10

Configuring Application Inspection

Getting Started with Application Layer Protocol Inspection

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path (see the "Stateful Inspection Overview" section on page 1-24 for more information about the fast path). As a result, inspection engines can affect overall throughput. Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- Information about Application Layer Protocol Inspection, page 57-1
- Guidelines and Limitations, page 57-3
- Default Settings, page 57-4
- Configuring Application Layer Protocol Inspection, page 57-6

Information about Application Layer Protocol Inspection

This section includes the following topics:

- How Inspection Engines Work, page 57-1
- When to Use Application Protocol Inspection, page 57-2

How Inspection Engines Work

As illustrated in Figure 57-1, the ASA uses three databases for its basic operation:

- Access lists—Used for authentication and authorization of connections based on specific networks, hosts, and services (TCP/UDP port numbers).
- Inspections—Contains a static, predefined set of application-level inspection functions.
- Connections (XLATE and CONN tables)—Maintains state and other information about each established connection. This information is used by the Adaptive Security Algorithm and cut-through proxy to efficiently forward traffic within established sessions.

Figure 57-1 How Inspection Engines Work

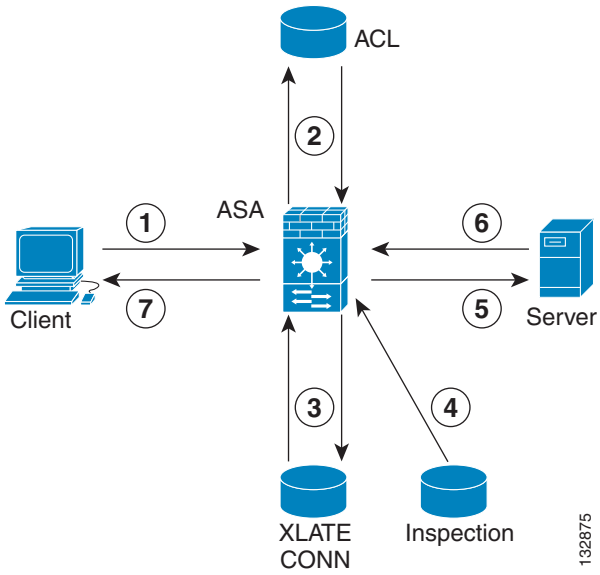

In Figure 57-1, operations are numbered in the order they occur, and are described as follows:

1. A TCP SYN packet arrives at the ASA to establish a new connection.
2. The ASA checks the access list database to determine if the connection is permitted.
3. The ASA creates a new entry in the connection database (XLATE and CONN tables).
4. The ASA checks the Inspections database to determine if the connection requires application-level inspection.
5. After the application inspection engine completes any required operations for the packet, the ASA forwards the packet to the destination system.
6. The destination system responds to the initial request.
7. The ASA receives the reply packet, looks up the connection in the connection database, and forwards the packet because it belongs to an established session.

The default configuration of the ASA includes a set of application inspection entries that associate supported protocols with specific TCP or UDP port numbers and that identify any special handling required.

When to Use Application Protocol Inspection

When a user establishes a connection, the ASA checks the packet against 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 ASA.

If you use applications like these, then you need to enable application inspection.

When you enable application inspection for a service that embeds IP addresses, the ASA translates embedded addresses and updates any checksum or other fields that are affected by the translation.

When you enable application inspection for a service that uses dynamically assigned ports, the ASA monitors sessions to identify the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

State information for multimedia sessions that require inspection are not passed over the state link for stateful failover. The exception is GTP, which is replicated over the state link.

IPv6 Guidelines

Supports IPv6 for the following inspections:

- DNS
- FTP
- HTTP
- ICMP
- SIP
- SMTP
- IPsec pass-through
- IPv6

Supports NAT64 for the following inspections:

- DNS
- FTP
- HTTP
- ICMP

Additional Guidelines and Limitations

Some inspection engines do not support PAT, NAT, outside NAT, or NAT between same security interfaces. See "Default Settings" for more information about NAT support.

For all the application inspections, the adaptive security appliance limits the number of simultaneous, active data connections to 200 connections. For example, if an FTP client opens multiple secondary connections, the FTP inspection engine allows only 200 active connections and the 201 connection is dropped and the adaptive security appliance generates a system error message.

Inspected protocols are subject to advanced TCP-state tracking, and the TCP state of these connections is not automatically replicated. While these connections are replicated to the standby unit, there is a best-effort attempt to re-establish a TCP state.

Default Settings

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 57-1 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 57-1 Supported Application Inspection Engines

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
CTIQBE	TCP/2748	No extended PAT. No NAT64.	—	—
DCERPC	TCP/135	No NAT64.	—	—
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	No extended PAT. No NAT64.	—	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. No extended PAT. No NAT64.	ITU-T H.323, H.245, H225.0, Q.931, Q.932	—
HTTP	TCP/80	—	RFC 2616	Beware of MTU limitations stripping ActiveX and Java. If the MTU is too small to allow the Java or ActiveX tag to be included in one packet, stripping may not occur.
ICMP	—	—	—	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 extended PAT. No NAT64.	—	—

Table 57-1 Supported Application Inspection Engines (continued)

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
Instant Messaging (IM)	Varies by client	No extended PAT. No NAT64.	RFC 3860	—
IP Options	—	No NAT64.	RFC 791, RFC 2113	All IP Options traffic is matched in the default class map.
MGCP	UDP/2427, 2727	No extended PAT. No NAT64.	RFC 2705bis-05	—
MMP	TCP 5443	No extended PAT. No NAT64.	—	—
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	—	NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
PPTP	TCP/1723	No NAT64.	RFC 2637	—
RADIUS Accounting	1646	No NAT64.	RFC 2865	—
RSH	TCP/514	No PAT. No NAT64.	Berkeley UNIX	—
RTSP	TCP/554	No extended PAT. No outside NAT. No NAT64.	RFC 2326, 2327, 1889	No handling for HTTP cloaking.
ScanSafe	TCP/80 TCP/413	—	—	These ports are not included in the default-inspection-traffic class for the ScanSafe inspection.
SIP	TCP/5060 UDP/5060	No outside NAT. No NAT on same security interfaces. No extended PAT. No NAT64.	RFC 2543	—
SKINNY (SCCP)	TCP/2000	No outside NAT. No NAT on same security interfaces. No extended PAT. No NAT64.	—	Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.
SMTP and **ESMTP**	TCP/25	No NAT64.	RFC 821, 1123	—
SNMP	UDP/161, 162	No NAT or PAT.	RFC 1155, 1157, 1212, 1213, 1215	v.2 RFC 1902-1908; v.3 RFC 2570-2580.
SQL*Net	TCP/1521	No extended PAT. No NAT64.	—	v.1 and v.2.

Table 57-1 Supported Application Inspection Engines (continued)

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
Sun RPC over UDP and TCP	UDP/111	No extended PAT. No NAT64.	—	The default rule includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new rule that matches TCP port 111 and performs Sun RPC inspection.
TFTP	UDP/69	No NAT64.	RFC 1350	Payload IP addresses are not translated.
WAAS	—	No extended PAT. No NAT64.	—	—
XDCMP	UDP/177	No extended PAT. No NAT64.	—	—

1. Inspection engines that are enabled by default for the default port are in bold.
2. The ASA is in compliance with these standards, but it does not enforce compliance on packets being inspected. For example, FTP commands are supposed to be in a particular order, but the ASA does not enforce the order.

Configuring Application Layer Protocol Inspection

This feature uses Security Policy Rules to create a service policy. Service policies provide a consistent and flexible way to configure ASA features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. See Chapter 44, "Configuring a Service Policy," for more information.

Inspection is enabled by default for some applications. See the "Default Settings" section for more information. Use this section to modify your inspection policy.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Service Policy Rules**.

Step 2 Add or edit a service policy rule according to the "Adding a Service Policy Rule for Through Traffic" section on page 44-8.

If you want to match non-standard ports, then create a new rule for the non-standard ports. See the "Default Settings" section on page 57-4 for the standard ports for each inspection engine. You can combine multiple rules in the same service policy if desired, so you can create one rule to match certain traffic, and another to match different traffic. However, if traffic matches a rule that contains an inspection action, and then matches another rule that also has an inspection action, only the first matching rule is used.

Step 3 In the Edit Service Policy Rule > Rule Actions dialog box, click the **Protocol Inspection** tab.

For a new rule, the dialog box is called Add Service Policy Rule Wizard - Rule Actions.

Step 4 Select each inspection type that you want to apply.

Step 5 (Optional) Some inspection engines let you control additional parameters when you apply the inspection to the traffic. Click **Configure** for each inspection type to configure an inspect map.

You can either choose an existing map, or create a new one. You can predefine inspect maps in the Configuration > Firewall > Objects > Inspect Maps pane.

Step 6 You can configure other features for this rule if desired using the other Rule Actions tabs.

Step 7 Click **OK** (or **Finish** from the wizard).

Configuring Inspection of Basic Internet Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- DNS Inspection, page 58-1
- FTP Inspection, page 58-12
- HTTP Inspection, page 58-21
- ICMP Inspection, page 58-34
- ICMP Error Inspection, page 58-34
- Instant Messaging Inspection, page 58-35
- IP Options Inspection, page 58-36
- IPsec Pass Through Inspection, page 58-40
- IPv6 Inspection, page 58-43
- NetBIOS Inspection, page 58-45
- PPTP Inspection, page 58-46
- SMTP and Extended SMTP Inspection, page 58-47
- TFTP Inspection, page 58-55

DNS Inspection

This section describes DNS application inspection. This section includes the following topics:

- How DNS Application Inspection Works, page 58-2
- How DNS Rewrite Works, page 58-3
- Configuring DNS Rewrite, page 58-3

- Select DNS Inspect Map, page 58-5
- DNS Class Map, page 58-6
- Add/Edit DNS Traffic Class Map, page 58-6
- Add/Edit DNS Match Criterion, page 58-7
- DNS Inspect Map, page 58-8
- Add/Edit DNS Policy Map (Security Level), page 58-9
- Add/Edit DNS Policy Map (Details), page 58-11

How DNS Application Inspection Works

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

When DNS inspection is enabled, which is the default, the ASA performs the following additional tasks:

- Translates the DNS record based on the configuration completed using the **alias**, **static** and **nat** commands (DNS Rewrite). 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 ASA performs reassembly as needed to verify that the packet length is less than the maximum length configured. The ASA drops the packet if it exceeds the maximum length.

Note If you enter the **inspect dns** command without the **maximum-length** option, DNS packet size is not checked

- 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 ASA 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 58-1, the DNS server resides on the external (ISP) network The real address of the server (192.168.100.1) has been mapped using the **static** command 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 ASA 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 ASA 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 58-1 Translating the Address in a DNS Reply (DNS Rewrite)

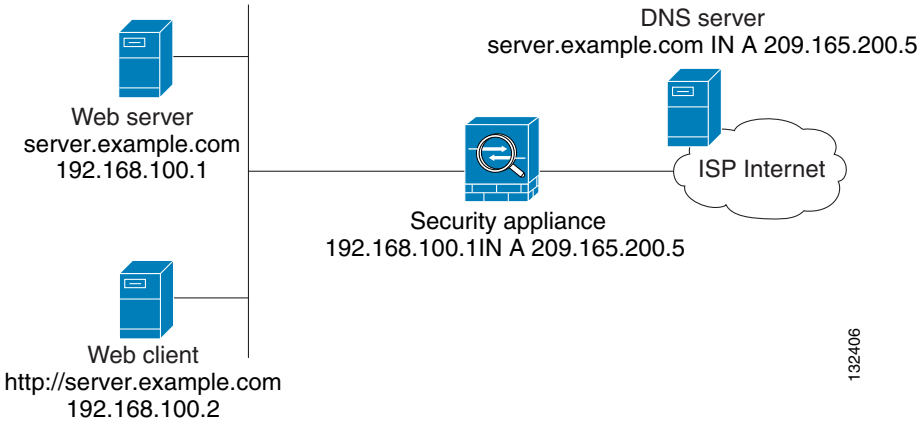

DNS rewrite also works if the client making the DNS request is on a DMZ network and the DNS server is on an inside interface.

Configuring DNS Rewrite

You configure DNS rewrite using the NAT configuration.

Figure 58-2 provides a more complex scenario to illustrate how DNS inspection allows NAT to operate transparently with a DNS server with minimal configuration.

Figure 58-2 DNS Rewrite with Three NAT Zones

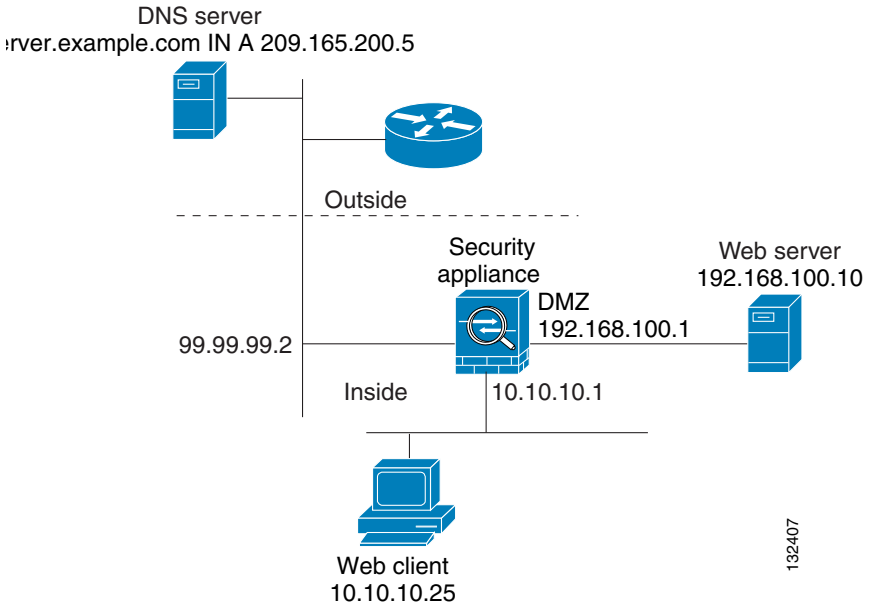

In Figure 58-2, a web server, server.example.com, has the real address 192.168.100.10 on the DMZ interface of the ASA. A web client with the IP address 10.10.10.25 is on the inside interface and a public DNS server is on the outside interface. The site NAT policies are as follows:

- The outside DNS server holds the authoritative address record for server.example.com.
- Hosts on the outside network can contact the web server with the domain name server.example.com through the outside DNS server or with the IP address 209.165.200.5.
- Clients on the inside network can access the web server with the domain name server.example.com through the outside DNS server or with the IP address 192.168.100.10.

When a host or client on any interface accesses the DMZ web server, it queries the public DNS server for the A-record of server.example.com. The DNS server returns the A-record showing that server.example.com binds to address 209.165.200.5.

When a web client on the *outside* network attempts to access http://server.example.com, the sequence of events is as follows:

1. The host running the web client sends the DNS server a request for the IP address of server.example.com.
2. The DNS server responds with the IP address 209.165.200.225 in the reply.
3. The web client sends its HTTP request to 209.165.200.225.
4. The packet from the outside host reaches the ASA at the outside interface.
5. The static rule translates the address 209.165.200.225 to 192.168.100.10 and the ASA directs the packet to the web server on the DMZ.

When a web client on the *inside* network attempts to access http://server.example.com, the sequence of events is as follows:

1. The host running the web client sends the DNS server a request for the IP address of server.example.com.
2. The DNS server responds with the IP address 209.165.200.225 in the reply.

3. The ASA receives the DNS reply and submits it to the DNS application inspection engine.
4. The DNS application inspection engine does the following:
   a. Searches for any NAT rule to undo the translation of the embedded A-record address "[outside]:209.165.200.5". In this example, it finds the following static configuration:

```
object network obj-192.168.100.10-01
  host 192.168.100.10
  nat (dmz,outside) static 209.165.200.5 dns
```

   b. Uses the static rule to rewrite the A-record as follows because the **dns** option is included:

```
[outside]:209.165.200.225 --> [dmz]:192.168.100.10
```

   Note If the **dns** option were not included with the **nat** command, DNS Rewrite would not be performed and other processing for the packet continues.

   c. Searches for any NAT to translate the web server address, [dmz]:192.168.100.10, when communicating with the inside web client.

      No NAT rule is applicable, so application inspection completes.

      If a NAT rule (nat or static) were applicable, the **dns** option must also be specified. If the **dns** option were not specified, the A-record rewrite in step b would be reverted and other processing for the packet continues.
5. The ASA sends the HTTP request to server.example.com on the DMZ interface.

Select DNS Inspect Map

The Select DNS Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select DNS Inspect 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.
- Enable Botnet traffic filter DNS snooping— 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. We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the ASA. For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.
- Add—Opens the Add Policy Map dialog box for the inspection.

DNS Class Map

The DNS Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > DNS

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 is accessible as follows:

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map

The Add/Edit DNS Traffic Class Map dialog box lets you define a DNS class map.

Fields

- Name—Enter the name of the DNS class map, up to 40 characters in length.
- Description—Enter the description of the DNS class map.
- Add—Adds a DNS class map.
- Edit—Edits a DNS class map.

- Delete—Deletes a DNS class map.

Add/Edit DNS Match Criterion

The Add/Edit DNS Match Criterion dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map > Add/Edit DNS Match Criterion

The Add/Edit DNS Match Criterion dialog box lets you define the match criterion and value for the DNS class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of DNS traffic to match.
  - Header Flag—Match a DNS flag in the header.
  - Type—Match a DNS query or resource record type.
  - Class—Match a DNS query or resource record class.
  - Question—Match a DNS question.
  - Resource Record—Match a DNS resource record.
  - Domain Name—Match a domain name from a DNS query or resource record.
- Header Flag Criterion Values—Specifies the value details for the DNS header flag match.
  - Match Option—Specifies either an exact match or match all bits (bit mask match).
  - Match Value—Specifies to match either the header flag name or the header flag value.

    Header Flag Name—Lets you select one or more header flag names to match, including AA (authoritative answer), QR (query), RA (recursion available), RD (recursion denied), TC (truncation) flag bits.

    Header Flag Value—Lets you enter an arbitrary 16-bit value in hex to match.
- Type Criterion Values—Specifies the value details for the DNS type match.
  - DNS Type Field Name—Lists the DNS types to select.

    A—IPv4 address

    NS—Authoritative name server

    CNAME—Canonical name

    SOA—Start of a zone of authority

    TSIG—Transaction signature

    IXFR—Incremental (zone) transfer

    AXFR—Full (zone) transfer
  - DNS Type Field Value—Specifies to match either a DNS type field value or a DNS type field range.

    Value—Lets you enter an arbitrary value between 0 and 65535 to match.

Range—Lets you enter a range match. Both values between 0 and 65535.

- Class Criterion Values—Specifies the value details for the DNS class match.
  - DNS Class Field Name—Specifies to match on internet, the DNS class field name.
  - DNS Class Field Value—Specifies to match either a DNS class field value or a DNS class field range.

    Value—Lets you enter an arbitrary value between 0 and 65535 to match.

    Range—Lets you enter a range match. Both values between 0 and 65535.
- Question Criterion Values—Specifies to match on the DNS question section.
- Resource Record Criterion Values—Specifies to match on the DNS resource record section.
  - Resource Record— Lists the sections to match.

    Additional—DNS additional resource record

    Answer—DNS answer resource record

    Authority—DNS authority resource record
- Domain Name Criterion Values—Specifies to match on the DNS domain name.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

DNS Inspect Map

The DNS Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > DNS

The DNS pane lets you view previously configured DNS application inspection maps. A DNS map lets you change the default configuration values used for DNS application inspection.

DNS application inspection supports DNS message controls that provide protection against DNS spoofing and cache poisoning. User configurable rules allow certain DNS types to be allowed, dropped, and/or logged, while others are blocked. Zone transfer can be restricted between servers with this function, for example.

The Recursion Desired and Recursion Available flags in the DNS header can be masked to protect a public server from attack if that server only supports a particular internal zone. In addition, DNS randomization can be enabled avoid spoofing and cache poisoning of servers that either do not support randomization, or utilize a weak pseudo random number generator. Limiting the domain names that can be queried also restricts the domain names which can be queried, which protects the public server further.

A configurable DNS mismatch alert can be used as notification if an excessive number of mismatching DNS responses are received, which could indicate a cache poisoning attack. In addition, a configurable check to enforce a Transaction Signature be attached to all DNS messages is also supported.

Fields

- 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.

Add/Edit DNS Policy Map (Security Level)

The Add/Edit DNS Policy Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > DNS > DNS Inspect Map > Basic View

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.

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.

FTP Inspection

This section describes the FTP inspection engine. This section includes the following topics:

- FTP Inspection Overview, page 58-12
- Using Strict FTP, page 58-13
- Select FTP Map, page 58-14
- FTP Class Map, page 58-14
- Add/Edit FTP Traffic Class Map, page 58-15
- Add/Edit FTP Match Criterion, page 58-15
- FTP Inspect Map, page 58-16

FTP Inspection Overview

The FTP application inspection inspects the FTP sessions and performs four tasks:

- Prepares dynamic secondary data connection
- Tracks the FTP command-response sequence
- Generates an audit trail
- Translates the embedded IP address

FTP application inspection prepares secondary channels for FTP data transfer. Ports for these channels are negotiated through PORT or PASV commands. The channels are allocated in response to a file upload, a file download, or a directory listing event.

Note If you disable FTP inspection engines with the **no inspect ftp** command, outbound users can start connections only in passive mode, and all inbound FTP is disabled.

Using Strict FTP

Using strict FTP increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests. To enable strict FTP, click the **Configure** button next to FTP on the Configuration > Firewall > Service Policy Rules > Edit Service Policy Rule > Rule Actions > Protocol Inspection tab.

After you enable the **strict** option on an interface, FTP inspection enforces the following behavior:

- An FTP command must be acknowledged before the ASA allows a new command.
- The ASA drops connections that send embedded commands.
- The 227 and PORT commands are checked to ensure they do not appear in an error string.

Caution Using the **strict** option may cause the failure of FTP clients that are not strictly compliant with FTP RFCs.

If the **strict** option is enabled, each FTP command and response sequence is tracked for the following anomalous activity:

- Truncated command—Number of commas in the PORT and PASV reply command is checked to see if it is five. If it is not five, then the PORT command is assumed to be truncated and the TCP connection is closed.
- Incorrect command—Checks the FTP command to see if it ends with <CR><LF> characters, as required by the RFC. If it does not, the connection is closed.
- Size of RETR and STOR commands—These are checked against a fixed constant. If the size is greater, then an error message is logged and the connection is closed.
- Command spoofing—The PORT command should always be sent from the client. The TCP connection is denied if a PORT command is sent from the server.
- Reply spoofing—PASV reply command (227) should always be sent from the server. The TCP connection is denied if a PASV reply command is sent from the client. This prevents the security hole when the user executes “227 xxxxx a1, a2, a3, a4, p1, p2.”
- TCP stream editing—The ASA closes the connection if it detects TCP stream editing.
- Invalid port negotiation—The negotiated dynamic port value is checked to see if it is less than 1024. As port numbers in the range from 1 to 1024 are reserved for well-known connections, if the negotiated port falls in this range, then the TCP connection is freed.
- Command pipelining—The number of characters present after the port numbers in the PORT and PASV reply command is cross checked with a constant value of 8. If it is more than 8, then the TCP connection is closed.
- The ASA replaces the FTP server response to the SYST command with a series of Xs. to prevent the server from revealing its system type to FTP clients. To override this default behavior, use the **no mask-syst-reply** command in the FTP map.

Select FTP Map

The Select FTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select FTP Map

The Select FTP Map dialog box lets you enable strict FTP application inspection, select an FTP map, or create a new FTP map. An FTP map lets you change the configuration values used for FTP application inspection.The Select FTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- FTP Strict (prevent web browsers from sending embedded commands in FTP requests)—Enables strict FTP application inspection, which causes the ASA to drop the connection when an embedded command is included in an FTP request.
- Use the default FTP inspection map—Specifies to use the default FTP map.
- Select an FTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

FTP Class Map

The FTP Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP

The FTP Class Map pane lets you configure FTP class maps for FTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the FTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the FTP class map.
  - Value—Shows the value to match in the FTP class map.
- Description—Shows the description of the class map.
- Add—Adds an FTP class map.
- Edit—Edits an FTP class map.
- Delete—Deletes an FTP class map.

Add/Edit FTP Traffic Class Map

The Add/Edit FTP Traffic Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP > Add/Edit FTP Traffic Class Map

The Add/Edit FTP Traffic Class Map dialog box lets you define a FTP class map.

Fields

- Name—Enter the name of the FTP class map, up to 40 characters in length.
- Description—Enter the description of the FTP class map.
- Add—Adds an FTP class map.
- Edit—Edits an FTP class map.
- Delete—Deletes an FTP class map.

Add/Edit FTP Match Criterion

The Add/Edit FTP Match Criterion dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP > Add/Edit FTP Traffic Class Map > Add/Edit FTP Match Criterion

The Add/Edit FTP Match Criterion dialog box lets you define the match criterion and value for the FTP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of FTP traffic to match.
  - Request-Command—Match an FTP request command.
  - File Name—Match a filename for FTP transfer.
  - File Type—Match a file type for FTP transfer.
  - Server—Match an FTP server.
  - User Name—Match an FTP user.
- Request-Command Criterion Values—Specifies the value details for the FTP request command match.
  - Request Command—Lets you select one or more request commands to match.

    APPE—Append to a file.

    CDUP—Change to the parent of the current directory.

    DELE—Delete a file at the server site.

    GET—FTP client command for the retr (retrieve a file) command.

    HELP—Help information from the server.

    MKD—Create a directory.

PUT—FTP client command for the stor (store a file) command.

RMD—Remove a directory.

RNFR—Rename from.

RNTO—Rename to.

SITE—Specify a server specific command.

STOU—Store a file with a unique name.

- File Name Criterion Values—Specifies to match on the FTP transfer filename.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- File Type Criterion Values—Specifies to match on the FTP transfer file type.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Server Criterion Values—Specifies to match on the FTP server.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- User Name Criterion Values—Specifies to match on the FTP user.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

FTP Inspect Map

The FTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP

The FTP pane lets you view previously configured FTP application inspection maps. An FTP map lets you change the default configuration values used for FTP application inspection.

FTP command filtering and security checks are provided using strict FTP inspection for improved security and control. Protocol conformance includes packet length checks, delimiters and packet format checks, command terminator checks, and command validation.

Blocking FTP based on user values is also supported so that it is possible for FTP sites to post files for download, but restrict access to certain users. You can block FTP connections based on file type, server name, and other attributes. System message logs are generated if an FTP connection is denied after inspection.

Fields

- FTP Inspect Maps—Table that lists the defined FTP inspect maps.
- Add—Configures a new FTP inspect map. To edit an FTP inspect map, choose the FTP entry in the FTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the FTP Inspect Maps table.
- Security Level—Select the security level (medium or low).
  - Low

    Mask Banner Disabled

    Mask Reply Disabled
  - Medium—Default.

    Mask Banner Enabled

    Mask Reply Enabled
  - File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
  - Customize—Opens the Add/Edit FTP Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.

File Type Filtering

The File Type Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > MIME File Type Filtering

The File Type Filtering dialog box lets you configure the settings for a file type filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add File Type Filter dialog box to add a file type filter.
- Edit—Opens the Edit File Type Filter dialog box to edit a file type filter.
- Delete—Deletes a file type filter.

- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit FTP Policy Map (Security Level)

The Add/Edit FTP Policy Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Basic View

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- Security Level—Select the security level (medium or low).
  - Low

    Mask Banner Disabled

    Mask Reply Disabled
  - Medium—Default.

    Mask Banner Enabled

    Mask Reply Enabled
  - File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
  - Default Level—Sets the security level back to the default level of Medium.
- Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit FTP Policy Map (Details)

The Add/Edit FTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Advanced View

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- Security Level—Shows the security level and file type filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the FTP inspect map.
  - Mask greeting banner from the server—Masks the greeting banner from the FTP server to prevent the client from discovering server information.

  - Mask reply to SYST command—Masks the reply to the syst command to prevent the client from discovering server information.
- Inspections—Tab that shows you the FTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the FTP inspection.
  - Value—Shows the value to match in the FTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add FTP Inspect dialog box to add an FTP inspection.
  - Edit—Opens the Edit FTP Inspect dialog box to edit an FTP inspection.
  - Delete—Deletes an FTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit FTP Map

The Add/Edit FTP Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Advanced View > Add/Edit FTP Inspect

The Add/Edit FTP Inspect dialog box lets you define the match criterion and value for the FTP inspect map.

Fields

- Single Match—Specifies that the FTP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of FTP traffic to match.
  - Request Command—Match an FTP request command.
  - File Name—Match a filename for FTP transfer.
  - File Type—Match a file type for FTP transfer.
  - Server—Match an FTP server.
  - User Name—Match an FTP user.
- Request Command Criterion Values—Specifies the value details for FTP request command match.
  - Request Command:

    APPE—Command that appends to a file.

    CDUP—Command that changes to the parent directory of the current working directory.

    DELE—Command that deletes a file.

    GET—Command that gets a file.

    HELP—Command that provides help information.

MKD—Command that creates a directory.

PUT—Command that sends a file.

RMD—Command that deletes a directory.

RNFR—Command that specifies rename-from filename.

RNTO—Command that specifies rename-to filename.

SITE—Commands that are specific to the server system. Usually used for remote administration.

STOU—Command that stores a file using a unique filename.

- File Name Criterion Values—Specifies the value details for FTP filename match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- File Type Criterion Values—Specifies the value details for FTP file type match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Server Criterion Values—Specifies the value details for FTP server match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- User Name Criterion Values—Specifies the value details for FTP user name match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Multiple Matches—Specifies multiple matches for the FTP inspection.
  - FTP Traffic Class—Specifies the FTP traffic class match.
  - Manage—Opens the Manage FTP Class Maps dialog box to add, edit, or delete FTP Class Maps.

- Action—Reset.
- Log—Enable or disable.

Verifying and Monitoring FTP Inspection

FTP application inspection generates the following log messages:

- An Audit record 303002 is generated for each file that is retrieved or uploaded.
- The FTP command is checked to see if it is RETR or STOR and the retrieve and store commands are logged.
- The username is obtained by looking up a table providing the IP address.
- The username, source IP address, destination IP address, NAT address, and the file operation are logged.
- Audit record 201005 is generated if the secondary dynamic channel preparation failed due to memory shortage.

In conjunction with NAT, the FTP application inspection translates the IP address within the application payload. This is described in detail in RFC 959.

HTTP Inspection

This section describes the HTTP inspection engine. This section includes the following topics:

- HTTP Inspection Overview, page 58-21
- Select HTTP Map, page 58-22
- HTTP Class Map, page 58-22
- Add/Edit HTTP Traffic Class Map, page 58-23
- Add/Edit HTTP Match Criterion, page 58-23
- HTTP Inspect Map, page 58-27
- “URI Filtering” section on page 58-28
- “Add/Edit HTTP Policy Map (Security Level)” section on page 58-28
- “Add/Edit HTTP Policy Map (Details)” section on page 58-29
- “Add/Edit HTTP Map” section on page 58-30

HTTP Inspection Overview

Use the HTTP inspection engine to protect against specific attacks and other threats that are associated with HTTP traffic. HTTP inspection performs several functions:

- Enhanced HTTP inspection
- URL screening through N2H2 or Websense

  See Information About URL Filtering, page 79-2 for information.
- 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 map, can help prevent attackers from using HTTP messages for circumventing network security policy. It verifies the following for all HTTP messages:

- Conformance to RFC 2616
- Use of RFC-defined methods only.
- Compliance with the additional criteria.

Select HTTP Map

The Select HTTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select HTTP Map

The Select HTTP Map dialog box lets you select or create a new HTTP map. An HTTP map lets you change the configuration values used for HTTP application inspection. The Select HTTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default HTTP inspection map—Specifies to use the default HTTP map.
- Select an HTTP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

HTTP Class Map

The HTTP Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP

The HTTP Class Map pane lets you configure HTTP class maps for HTTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the HTTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the HTTP class map.
  - Value—Shows the value to match in the HTTP class map.
- Description—Shows the description of the class map.
- Add—Adds an HTTP class map.
- Edit—Edits an HTTP class map.

- Delete—Deletes an HTTP class map.

Add/Edit HTTP Traffic Class Map

The Add/Edit HTTP Traffic Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP > Add/Edit HTTP Traffic Class Map

The Add/Edit HTTP Traffic Class Map dialog box lets you define a HTTP class map.

Fields

- Name—Enter the name of the HTTP class map, up to 40 characters in length.
- Description—Enter the description of the HTTP class map.
- Add—Adds an HTTP class map.
- Edit—Edits an HTTP class map.
- Delete—Deletes an HTTP class map.

Add/Edit HTTP Match Criterion

The Add/Edit HTTP Match Criterion dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP > Add/Edit HTTP Traffic Class Map > Add/Edit HTTP Match Criterion

The Add/Edit HTTP Match Criterion dialog box lets you define the match criterion and value for the HTTP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.

- Criterion—Specifies which criterion of HTTP traffic to match.
  - Request/Response Content Type Mismatch—Specifies that the content type in the response must match one of the MIME types in the accept field of the request.
  - Request Arguments—Applies the regular expression match to the arguments of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
  - Request Body Length—Applies the regular expression match to the body of the request with field length greater than the bytes specified.

    Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

- Request Body—Applies the regular expression match to the body of the request.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Header Field Count—Applies the regular expression match to the header of the request with a maximum number of header fields.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Count—Enter the maximum number of header fields.

- Request Header Field Length—Applies the regular expression match to the header of the request with field length greater than the bytes specified.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

- Request Header Field—Applies the regular expression match to the header of the request.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Header Count—Applies the regular expression match to the header of the request with a maximum number of headers.

  Greater Than Count—Enter the maximum number of headers.

- Request Header Length—Applies the regular expression match to the header of the request with length greater than the bytes specified.

  Greater Than Length—Enter a header length value in bytes.

- Request Header non-ASCII—Matches non-ASCII characters in the header of the request.

- Request Method—Applies the regular expression match to the method of the request.

  Method—Specifies to match on a request method: bcopy, bdelete, bmove, bpropfind, bproppatch, connect, copy, delete, edit, get, getattribute, getattributenames, getproperties, head, index, lock, mkcol, mkdir, move, notify, options, poll, post, propfind, proppatch, put, revadd, revlabel, revlog, revnum, save, search, setattribute, startrev, stoprev, subscribe, trace, unedit, unlock, unsubscribe.

  Regular Expression—Specifies to match on a regular expression.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request URI Length—Applies the regular expression match to the URI of the request with length greater than the bytes specified.

  Greater Than Length—Enter a URI length value in bytes.

- Request URI—Applies the regular expression match to the URI of the request.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Response Body—Applies the regex match to the body of the response.

  ActiveX—Specifies to match on ActiveX.

  Java Applet—Specifies to match on a Java Applet.

  Regular Expression—Specifies to match on a regular expression.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Response Body Length—Applies the regular expression match to the body of the response with field length greater than the bytes specified.

Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

- Response Header Field Count—Applies the regular expression match to the header of the response with a maximum number of header fields.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Count—Enter the maximum number of header fields.

- Response Header Field Length—Applies the regular expression match to the header of the response with field length greater than the bytes specified.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

- Response Header Field—Applies the regular expression match to the header of the response.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Response Header Count—Applies the regular expression match to the header of the response with a maximum number of headers.

  Greater Than Count—Enter the maximum number of headers.

- Response Header Length—Applies the regular expression match to the header of the response with length greater than the bytes specified.

  Greater Than Length—Enter a header length value in bytes.

- Response Header non-ASCII—Matches non-ASCII characters in the header of the response.

- Response Status Line—Applies the regular expression match to the status line.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

HTTP Inspect Map

The HTTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP

The HTTP pane lets you view previously configured HTTP application inspection maps. An HTTP map lets you change the default configuration values used for HTTP application inspection.

HTTP application inspection scans HTTP headers and body, and performs various checks on the data. These checks prevent various HTTP constructs, content types, and tunneling and messaging protocols from traversing the security appliance.

HTTP application inspection can block tunneled applications and non-ASCII characters in HTTP requests and responses, preventing malicious content from reaching the web server. Size limiting of various elements in HTTP request and response headers, URL blocking, and HTTP server header type spoofing are also supported.

Fields

- HTTP Inspect Maps—Table that lists the defined HTTP inspect maps.
- Add—Configures a new HTTP inspect map. To edit an HTTP inspect map, choose the HTTP entry in the HTTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the HTTP Inspect Maps table.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Disabled

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - Medium

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Allow only GET, HEAD, and POST

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - High

Protocol violation action: Drop connection and log

Drop connections for unsafe methods: Allow only GET and HEAD.

Drop connections for requests with non-ASCII headers: Enabled

URI filtering: Not configured

Advanced inspections: Not configured

- URI Filtering—Opens the URI Filtering dialog box to configure URI filters.
- Customize—Opens the Edit HTTP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.

URI Filtering

The URI Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > URI Filtering

The URI Filtering dialog box lets you configure the settings for an URI filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add URI Filtering dialog box to add a URI filter.
- Edit—Opens the Edit URI Filtering dialog box to edit a URI filter.
- Delete—Deletes an URI filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit HTTP Policy Map (Security Level)

The Add/Edit HTTP Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Basic View

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

Protocol violation action: Drop connection

Drop connections for unsafe methods: Disabled

Drop connections for requests with non-ASCII headers: Disabled

URI filtering: Not configured

Advanced inspections: Not configured

- Medium

  Protocol violation action: Drop connection

  Drop connections for unsafe methods: Allow only GET, HEAD, and POST

  Drop connections for requests with non-ASCII headers: Disabled

  URI filtering: Not configured

  Advanced inspections: Not configured

- High

  Protocol violation action: Drop connection and log

  Drop connections for unsafe methods: Allow only GET and HEAD.

  Drop connections for requests with non-ASCII headers: Enabled

  URI filtering: Not configured

  Advanced inspections: Not configured

- URI Filtering—Opens the URI Filtering dialog box which lets you configure the settings for an URI filter.
- Default Level—Sets the security level back to the default.

- Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit HTTP Policy Map (Details)

The Add/Edit HTTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Advanced View

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Shows the security level and URI filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the HTTP inspect map.
  - Check for protocol violations—Checks for HTTP protocol violations.

    Action—Drop Connection, Reset, Log.

    Log—Enable or disable.
  - Spoof server string—Replaces the server HTTP header value with the specified string.

Spoof String—Enter a string to substitute for the server header field. Maximum is 82 characters.

- Body Match Maximum—The maximum number of characters in the body of an HTTP message that should be searched in a body match. Default is 200 bytes. A large number will have a significant impact on performance.

- Inspections—Tab that shows you the HTTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the HTTP inspection.
  - Value—Shows the value to match in the HTTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add HTTP Inspect dialog box to add an HTTP inspection.
  - Edit—Opens the Edit HTTP Inspect dialog box to edit an HTTP inspection.
  - Delete—Deletes an HTTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit HTTP Map

The Add/Edit HTTP Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Advanced View > Add/Edit HTTP Inspect

The Add/Edit HTTP Inspect dialog box lets you define the match criterion and value for the HTTP inspect map.

Fields

- Single Match—Specifies that the HTTP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of HTTP traffic to match.
  - Request/Response Content Type Mismatch—Specifies that the content type in the response must match one of the MIME types in the accept field of the request.
  - Request Arguments—Applies the regular expression match to the arguments of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
  - Request Body Length—Applies the regular expression match to the body of the request with field length greater than the bytes specified.

Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

- Request Body—Applies the regular expression match to the body of the request.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Header Field Count—Applies the regular expression match to the header of the request with a maximum number of header fields.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Count—Enter the maximum number of header fields.

- Request Header Field Length—Applies the regular expression match to the header of the request with field length greater than the bytes specified.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

- Request Header Field—Applies the regular expression match to the header of the request.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Header Count—Applies the regular expression match to the header of the request with a maximum number of headers.

  Greater Than Count—Enter the maximum number of headers.

- Request Header Length—Applies the regular expression match to the header of the request with length greater than the bytes specified.

  Greater Than Length—Enter a header length value in bytes.

- Request Header non-ASCII—Matches non-ASCII characters in the header of the request.

- Request Method—Applies the regular expression match to the method of the request.

  Method—Specifies to match on a request method: bcopy, bdelete, bmove, bpropfind, bproppatch, connect, copy, delete, edit, get, getattribute, getattributenames, getproperties, head, index, lock, mkcol, mkdir, move, notify, options, poll, post, propfind, proppatch, put, revadd, revlabel, revlog, revnum, save, search, setattribute, startrev, stoprev, subscribe, trace, unedit, unlock, unsubscribe.

  Regular Expression—Specifies to match on a regular expression.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request URI Length—Applies the regular expression match to the URI of the request with length greater than the bytes specified.

  Greater Than Length—Enter a URI length value in bytes.

- Request URI—Applies the regular expression match to the URI of the request.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Response Body—Applies the regex match to the body of the response.

  ActiveX—Specifies to match on ActiveX.

  Java Applet—Specifies to match on a Java Applet.

  Regular Expression—Specifies to match on a regular expression.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Response Body Length—Applies the regular expression match to the body of the response with field length greater than the bytes specified.

  Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

– Response Header Field Count—Applies the regular expression match to the header of the response with a maximum number of header fields.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Count—Enter the maximum number of header fields.

– Response Header Field Length—Applies the regular expression match to the header of the response with field length greater than the bytes specified.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

– Response Header Field—Applies the regular expression match to the header of the response.

  Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Response Header Count—Applies the regular expression match to the header of the response with a maximum number of headers.

  Greater Than Count—Enter the maximum number of headers.

– Response Header Length—Applies the regular expression match to the header of the response with length greater than the bytes specified.

Greater Than Length—Enter a header length value in bytes.

- Response Header non-ASCII—Matches non-ASCII characters in the header of the response.
- Response Status Line—Applies the regular expression match to the status line.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Multiple Matches—Specifies multiple matches for the HTTP inspection.
  - H323 Traffic Class—Specifies the HTTP traffic class match.
  - Manage—Opens the Manage HTTP Class Maps dialog box to add, edit, or delete HTTP Class Maps.
- Action—Drop connection, reset, or log.
- Log—Enable or disable.

ICMP Inspection

The ICMP inspection engine allows ICMP traffic to have a "session" so it can be inspected like TCP and UDP traffic. Without the ICMP inspection engine, we recommend that you do not allow ICMP through the ASA in an 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 ASA creates translation sessions for intermediate hops that send ICMP error messages, based on the NAT configuration. The ASA overwrites the packet with the translated IP addresses.

When disabled, the ASA does not create translation sessions for intermediate nodes that generate ICMP error messages. ICMP error messages generated by the intermediate nodes between the inside host and the ASA reach the outside host without consuming any additional NAT resource. This is undesirable when an outside host uses the traceroute command to trace the hops to the destination on the inside of the ASA. When the ASA does not translate the intermediate hops, all the intermediate hops appear with the mapped destination IP address.

The ICMP payload is scanned to retrieve the five-tuple from the original packet. Using the retrieved five-tuple, a lookup is performed to determine the original address of the client. The ICMP error inspection engine makes the following changes to the ICMP packet:

- In the IP Header, the mapped IP is changed to the real IP (Destination Address) and the IP checksum is modified.
- In the ICMP Header, the ICMP checksum is modified due to the changes in the ICMP packet.
- In the Payload, the following changes are made:
  - Original packet mapped IP is changed to the real IP

- Original packet mapped port is changed to the real Port
- Original packet IP checksum is recalculated

Instant Messaging Inspection

This section describes the IM inspection engine. This section includes the following topics:

- IM Inspection Overview, page 58-35
- Select IM Map, page 58-36

IM Inspection Overview

The IM inspect engine lets you apply fine grained controls on the IM application to control the network usage and stop leakage of confidential data, propagation of worms, and other threats to the corporate network.

Adding a Class Map for IM Inspection

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure IP Options inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1 Choose **Configuration > Firewall > Objects > Class Maps > Instant Messaging (IM)**. The table displaying the configured class maps for Instant Messaging Inspection appears.

Step 2 To add a new class map, click **Add**. The Add Instant Messaging (IM) Traffic Class Map dialog box appears.

Step 3 Enter a name for the class map.

Step 4 (Optional) Enter a description for the class map. The description can contain up to 200 characters.

Step 5 In the Match Option field, click an option for the class map:

- Match All—Specifies that traffic must match all criteria to match the class map. By default, the Match All option is selected.
- Match Any—Specifies that the traffic matches the class map if it matches at least one of the criteria.

Step 6 Click **Add** to add a match criteria for the class map. The Add Instant Messaging (IM) Match Criterion dialog box appears.

Step 7 In the Match Type field, click the Match or No Match radio button.

Step 8 In the Criterion drop-down list, select one of the following options and specify the criteria value. Depending on which option you select, the Value fields dynamically refresh to display the appropriate values for that criteria.

- Protocol—Select to match traffic of a specific IM protocol, such as Yahoo Messenger or MSN Messenger.
- Service—Select to match a specific IM service, such as chat, file-transfer, webcam, voice-chat, conference, or games.

- Version—Select to match the version of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

  See Configuring Regular Expressions, page 25-10.

- Client Login Name—Select to match the source login name of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

  See Configuring Regular Expressions, page 25-10.

- Client Peer Login Name—Select to match the destination login name of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

  See Configuring Regular Expressions, page 25-10.

- Source IP Address—Select to match the source IP address of the IM message. In the Value fields, enter the IP address and netmask of the message source.

- Destination IP Address—Select to match the destination IP address of the IM message. In the Value fields, enter the IP address and netmask of the message destination.

- Filename—Select to match the filename of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

  See Configuring Regular Expressions, page 25-10.

Step 9 Click **OK** to save the criteria. The Add Instant Messaging (IM) Match Criterion dialog box closes and the criteria appears in the Match Criterion table.

Step 10 Click **OK** to save the class map.

Select IM Map

The Select IM Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IM Map

The Select IM Map dialog box lets you select or create a new IM map. An IM map lets you change the configuration values used for IM application inspection. The Select IM Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Add—Opens the Add Policy Map dialog box for the inspection.

IP Options Inspection

This section describes the IP Options inspection engine. This section includes the following topics:

- IP Options Inspection Overview, page 58-37
- Configuring IP Options Inspection, page 58-37
- Select IP Options Inspect Map, page 58-39
- IP Options Inspect Map, page 58-39
- Add/Edit IP Options Inspect Map, page 58-39

IP Options Inspection Overview

Each IP packet contains an IP header with the Options field. The Options field, commonly referred to as IP Options, provide for control functions that are required in some situations but unnecessary for most common communications. In particular, IP Options include provisions for time stamps, security, and special routing. Use of IP Options is optional, and the field can contain zero, one, or more options.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the ASA. Configuring this inspection instructs the ASA to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

IP Options inspection can check for the following three IP options in a packet:

- End of Options List (EOOL) or IP Option 0—This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.
- No Operation (NOP) or IP Option 1—The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.
- Router Alert (RTRALT) or IP Option 20—This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

Note IP Options inspection is included by default in the global inspection policy. Therefore, the ASA allows RSVP traffic that contains packets with the Router Alert option (option 20) when the ASA is in routed mode.

Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations.

When you configure the ASA to clear the Router Alert option from IP headers, the IP header changes in the following ways:

- The Options field is padded so that the field ends on a 32 bit boundary.
- Internet header length (IHL) changes.
- The total length of the packet changes.
- The checksum is recomputed.

If an IP header contains additional options other than EOOL, NOP, or RTRALT, regardless of whether the ASA is configured to allow these options, the ASA will drop the packet.

Configuring IP Options Inspection

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure IP Options inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1 Open the Add Service Policy Rule Wizard by selecting **Configuration > Firewall > Service Policy Rules > Add**.

Perform the steps to complete the Service Policy, Traffic Classification Criteria, and Traffic Match - Destination Port pages of the wizard. See the "Adding a Service Policy Rule for Through Traffic" section on page 44-8.

The Add Service Policy Rule Wizard - Rule Actions dialog box opens.

Step 2 Check the **IP-Options** check box.

Step 3 Click **Configure**.

The Select IP Options Inspect Map dialog box opens.

Step 4 Perform one of the following:

- Click the **Use the default IP-Options inspection map** radio button to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).
- Click the **Select an IP-Options inspect map for fine control over inspection** radio button to select a defined application inspection map.
- Click Add to open the Add IP-Options Inspect Map dialog box and create a new inspection map.

Step 5 (Optional) If you clicked **Add** to create a new inspection map, define the following values for IP Options Inspection:

a. Enter a name for the inspection map.

b. Enter a description for the inspection map, up to 200 characters long.

c. From the Parameters area, select which IP options you want to pass through the ASA or clear and then pass through the ASA:

– Allow packets with the End of Options List (EOOL) option

This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.

– Allow packets with the No Operation (NOP) option

The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.

– Allow packets with the Router Alert (RTRALT) option

This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

– Clear the option value from the packets

When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the ASA.

d. Click **OK**.

Step 6 Click **OK**.

Step 7 Click **Finish**.

Select IP Options Inspect Map

The Select IP Options Inspect Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IM Map

The Select IP-Options Inspect Map dialog box lets you select or create a new IP Options inspection map. Use this inspection map to control whether the ASA drops, passes, or clears IP packets containing the following IP options—End of Options List, No Operations, and Router Alert.

Fields

- Use the default IP-Options inspection map—Specifies to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).
- Select an IP-Options map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add IP Options Inspect Map dialog box for the inspection.

IP Options Inspect Map

The IP Options Inspect Maps pane lets you view previously configured IP Options inspection maps. An IP Options inspection map lets you change the default configuration values used for IP Option inspection.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the security appliance. Configuring this inspection instructs the security appliance to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

In particular, you can control whether the security appliance drops, clears, or passes packets containing the Router Alert (RTRALT) option. Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations. Therefore, you can create IP Options inspection maps to pass packets containing the RTRALT option.

Fields

IP Options Inspect Maps—Table that lists the defined IP Options inspect maps.

Add—Configures a new IP Options inspect map.

Edit—Edits an existing IP Options inspect map. To edit an IP Options inspect map, choose the entry in the table and click Edit.

Delete—Deletes the inspect map selected in the IP Options Inspect Maps table.

Add/Edit IP Options Inspect Map

The Add/Edit IP Options Inspect Map lets you configure the settings for IP Options inspection maps.

Fields

- Name—When adding an IP Options inspection map, enter the name of the map. When editing a map, the name of the previously configured map is shown.
- Description—Enter the description of the IP Options inspection map, up to 200 characters in length.

- Parameters—Select which IP options you want to pass through the ASA or clear and then pass through the ASA:
  - Allow packets with the End of Options List (EOOL) option

  This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.
  - Allow packets with the No Operation (NOP) option

  The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.
  - Allow packets with the Router Alert (RTRALT) option

  This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.
  - Clear the option value from the packets

  When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the ASA.

IPsec Pass Through Inspection

This section describes the IPsec Pass Through inspection engine. This section includes the following topics:

- IPsec Pass Through Inspection Overview, page 58-40
- Select IPsec-Pass-Thru Map, page 58-41
- IPsec Pass Through Inspect Map, page 58-41
- Add/Edit IPsec Pass Thru Policy Map (Security Level), page 58-42
- Add/Edit IPsec Pass Thru Policy Map (Details), page 58-42

IPsec Pass Through Inspection Overview

Internet Protocol Security (IPsec) is a protocol suite for securing IP communications by authenticating and encrypting each IP packet of a data stream. IPsec also includes protocols for establishing mutual authentication between agents at the beginning of the session and negotiation of cryptographic keys to be used during the session. IPsec can be used to protect data flows between a pair of hosts (for example, computer users or servers), between a pair of security gateways (such as routers or firewalls), or between a security gateway and a host.

IPsec Pass Through application inspection provides convenient traversal of ESP (IP protocol 50) and AH (IP protocol 51) traffic associated with an IKE UDP port 500 connection. It avoids lengthy access list configuration to permit ESP and AH traffic and also provides security using timeout and max connections.

Specify IPsec Pass Through inspection parameters to identify a specific map to use for defining the parameters for the inspection. Configure a policy map for Specify IPsec Pass Through inspection to access the parameters configuration, which lets you specify the restrictions for ESP or AH traffic. You can set the per client max connections and the idle timeout in parameters configuration.

NAT and non-NAT traffic is permitted. However, PAT is not supported.

Select IPsec-Pass-Thru Map

The Select IPsec-Pass-Thru Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IPsec-Pass-Thru Map

The Select IPsec-Pass-Thru dialog box lets you select or create a new IPsec map. An IPsec map lets you change the configuration values used for IPsec application inspection. The Select IPsec Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IPsec inspection map—Specifies to use the default IPsec map.
- Select an IPsec map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

IPsec Pass Through Inspect Map

The IPsec Pass Through Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > IPsec Pass Through

The IPsec Pass Through pane lets you view previously configured IPsec Pass Through application inspection maps. An IPsec Pass Through map lets you change the default configuration values used for IPsec Pass Through application inspection. You can use an IPsec Pass Through map to permit certain flows without using an 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.

Add/Edit IPsec Pass Thru Policy Map (Security Level)

The Add/Edit IPsec Pass Thru Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > IPsec Pass Through>
IPsec Pass Through Inspect Map > Basic View

The Add/Edit IPsec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPsec Pass Thru application inspection maps.

Fields

- Name—When adding an IPsec Pass Thru map, enter the name of the IPsec Pass Thru map. When editing an IPsec Pass Thru map, the name of the previously configured IPsec Pass Thru map is shown.
- Security Level—Select the security level (high or low).
  - Low—Default.

    Maximum ESP flows per client: Unlimited.

    ESP idle timeout: 00:10:00.

    Maximum AH flows per client: Unlimited.

    AH idle timeout: 00:10:00.
  - High

    Maximum ESP flows per client:10.

    ESP idle timeout: 00:00:30.

    Maximum AH flows per client: 10.

    AH idle timeout: 00:00:30.
  - Default Level—Sets the security level back to the default level of Low.
- Details—Shows additional parameter settings to configure.

Add/Edit IPsec Pass Thru Policy Map (Details)

The Add/Edit IPsec Pass Thru Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps >IPsec Pass Through >
IPsec Pass Through Inspect Map > Advanced View

The Add/Edit IPsec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPsec Pass Thru application inspection maps.

Fields

- Name—When adding an IPsec Pass Thru map, enter the name of the IPsec Pass Thru map. When editing an IPsec Pass Thru map, the name of the previously configured IPsec Pass Thru map is shown.
- Description—Enter the description of the IPsec Pass Through map, up to 200 characters in length.
- Security Level—Shows the security level settings to configure.
- Parameters—Configures ESP and AH parameter settings.
  - Limit ESP flows per client—Limits ESP flows per client.

    Maximum—Specify maximum limit.
  - Apply ESP idle timeout—Applies ESP idle timeout.

    Timeout—Specify timeout.
  - Limit AH flows per client—Limits AH flows per client.

    Maximum—Specify maximum limit.
  - Apply AH idle timeout—Applies AH idle timeout.

    Timeout—Specify timeout.

IPv6 Inspection

You can configure IPv6 Inspection by using MPF rules to selectively block IPv6 traffic based on the extension header. IPv6 packets are subjected to an early security check. The ASA always passes hop-by-hop and destination option types of extension headers while blocking router header and no next header.

You can enable default IPv6 inspection or define IPv6 inspection. By defining an MPF policy map for IPv6 inspection you can configure the ASA to selectively drop IPv6 packets based on following types of extension headers found anywhere in the IPv6 packet:

- Hop-by-Hop Options
- Routing (Type 0)
- Fragment
- Destination Options
- Authentication
- Encapsulating Security Payload

In addition, default IPv6 inspection checks conformance to RFC 2460 for type and order of extension headers in IPv6 packets:

- IPv6 header
- Hop-by-Hop Options header (0)
- Destination Options header (60)
- Routing header (43)
- Fragment header (44)
- Authentication (51)
- Encapsulating Security Payload header(50)

- Destination Options header (60)
- No Next Header (59)

When a policy map is not configured for IPv6inspection or a configured policy map is not associated with an interface, the ASA drops packets with any mobility type and a routing-type IPv6 extension header that arrive at the interface.

When an IPv6 inspection policy map is created, the ASA automatically generates a configuration to drop packets that match header routing-type in the range 0-255.

Configuring an IPv6 Inspection Policy Map

You can configure a policy map for IPv6 inspection to handle IPv6 extension headers. The IPv6 policy map is applied to each classified IPv6 packet on the specified direction. Currently, only incoming IPv6 traffic is inspected.

Step 1 Choose **Configuration > Firewall > Objects > Inspect Maps > IPv6**. The Configure IPv6 Maps pane appears.

Step 2 Click **Add**. The Add IPv6 Inspection Map dialog box appears.

Step 3 Enter a name and description for the inspection map.

By default, the Enforcement tab is selected and the following options are selected:

- Permit only known extension headers
- Enforce extension header order

When **Permit only known extension headers** is selected, the ASA verifies the IPv6 extension header.

When **Enforce extension header order** is selected, the order of IPv6 extension headers as defined in the RFC 2460 Specification is enforced.

When these options are specified and an error is detected, the ASA drops the packet and logs the action.

Step 4 To configuring matching in the extension header, click the **Header Matches** tab.

Step 5 Click **Add** to add a match. The Add IPv6 Inspect dialog box appears.

Step 6 Select a criterion for the match.

When you select any of the following criteria, you can configure to the ASA to drop or log when an IPv6 packet arrives mathcing the criterion:

- Authentication (AH) header
- Destination Options header
- Encapsulating Security Payload (ESP) header
- Fragment header
- Hop-by-Hop Options header

When Routing header is selected and an IPv6 routing extension header is detected, the ASA takes the specified action when the routing type is matched or a number when the specified routing type range is matched.

When Header count is selected and an IPv6 routing extension header is detected, the ASA takes the specified action when number of IPv6 extension headers in the packet is more than the specified value.

When Routing header address count is selected, and an IPv6 routing extension header is detected, the ASA takes the specified action when the number of addresses in the type 0 routing header is more than the value you configure.

Step 7 Click **OK** to save the match criterion.

Step 8 Click **OK** to save the IPv6 inspect map.

NetBIOS Inspection

This section describes the IM inspection engine. This section includes the following topics:

- NetBIOS Inspection Overview, page 58-45
- Select NETBIOS Map, page 58-45
- "NetBIOS Inspect Map" section on page 58-45
- "Add/Edit NetBIOS Policy Map" section on page 58-46

NetBIOS Inspection Overview

NetBIOS inspection is enabled by default. The NetBios inspection engine translates IP addresses in the NetBios name service (NBNS) packets according to the ASA NAT configuration.

Select NETBIOS Map

The Select NETBIOS Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select NetBIOS Map

The Select NETBIOS Map dialog box lets you select or create a new NetBIOS map. A NetBIOS map lets you change the configuration values used for NetBIOS application inspection. The Select NetBIOS Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IM inspection map—Specifies to use the default NetBIOS map.
- Select a NetBIOS map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

NetBIOS Inspect Map

The NetBIOS Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > NetBIOS

The NetBIOS pane lets you view previously configured NetBIOS application inspection maps. A NetBIOS map lets you change the default configuration values used for NetBIOS application inspection.

NetBIOS application inspection performs NAT for the embedded IP address in the NetBIOS name service packets and NetBIOS datagram services packets. It also enforces protocol conformance, checking the various count and length fields for consistency.

Fields

- NetBIOS Inspect Maps—Table that lists the defined NetBIOS inspect maps.
- Add—Configures a new NetBIOS inspect map.
- Edit—Edits the selected NetBIOS entry in the NetBIOS Inspect Maps table.
- Delete—Deletes the inspect map selected in the NetBIOS Inspect Maps table.

Add/Edit NetBIOS Policy Map

The Add/Edit NetBIOS Policy Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > NetBIOS > NetBIOS Inspect Map > View

The Add/Edit NetBIOS Policy Map pane lets you configure the protocol violation settings for NetBIOS application inspection maps.

Fields

- Name—When adding a NetBIOS map, enter the name of the NetBIOS map. When editing an NetBIOS map, the name of the previously configured NetBIOS map is shown.
- Description—Enter the description of the NetBIOS map, up to 200 characters in length.
- Check for protocol violations—Checks for protocol violations and executes specified action.
  - Action—Drop packet or log.
  - Log—Enable or disable.

PPTP Inspection

PPTP is a protocol for tunneling PPP traffic. A PPTP session is composed of one TCP channel and usually two PPTP GRE tunnels. The TCP channel is the control channel used for negotiating and managing the PPTP GRE tunnels. The GRE tunnels carries PPP sessions between the two hosts.

When enabled, PPTP application inspection inspects PPTP protocol packets and dynamically creates the GRE connections and xlates necessary to permit PPTP traffic. Only Version 1, as defined in RFC 2637, is supported.

PAT is only performed for the modified version of GRE [RFC 2637] when negotiated over the PPTP TCP control channel. Port Address Translation is *not* performed for the unmodified version of GRE [RFC 1701, RFC 1702].

Specifically, the ASA inspects the PPTP version announcements and the outgoing call request/response sequence. Only PPTP Version 1, as defined in RFC 2637, is inspected. Further inspection on the TCP control channel is disabled if the version announced by either side is not Version 1. In addition, the outgoing-call request and reply sequence are tracked. Connections and xlates are dynamic allocated as necessary to permit subsequent secondary GRE data traffic.

The PPTP inspection engine must be enabled for PPTP traffic to be translated by PAT. Additionally, PAT is only performed for a modified version of GRE (RFC2637) and only if it is negotiated over the PPTP TCP control channel. PAT is not performed for the unmodified version of GRE (RFC 1701 and RFC 1702).

As described in RFC 2637, the PPTP protocol is mainly used for the tunneling of PPP sessions initiated from a modem bank PAC (PPTP Access Concentrator) to the headend PNS (PPTP Network Server). When used this way, the PAC is the remote client and the PNS is the server.

However, when used for VPN by Windows, the interaction is inverted. The PNS is a remote single-user PC that initiates connection to the head-end PAC to gain access to a central network.

SMTP and Extended SMTP Inspection

This section describes the IM inspection engine. This section includes the following topics:

- SMTP and ESMTP Inspection Overview, page 58-47
- Select ESMTP Map, page 58-48
- ESMTP Inspect Map, page 58-48
- MIME File Type Filtering, page 58-49
- Add/Edit ESMTP Policy Map (Security Level), page 58-50
- Add/Edit ESMTP Policy Map (Details), page 58-51
- Add/Edit ESMTP Inspect, page 58-52

SMTP and ESMTP Inspection Overview

ESMTP application inspection provides improved protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the ASA and by adding monitoring capabilities.

ESMTP is an enhancement to the SMTP protocol and is similar is most respects to SMTP. For convenience, the term SMTP is used in this document to refer to both SMTP and ESMTP. The application inspection process for extended SMTP is similar to SMTP application inspection and includes support for SMTP sessions. Most commands used in an extended SMTP session are the same as those used in an SMTP session but an ESMTP session is considerably faster and offers more options related to reliability and security, such as delivery status notification.

Extended SMTP application inspection adds support for these extended SMTP commands, including AUTH, EHLO, ETRN, HELP, SAML, SEND, SOML, STARTTLS, and VRFY. Along with the support for seven RFC 821 commands (DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET), the ASA supports a total of fifteen SMTP commands.

Other extended SMTP commands, such as ATRN, ONEX, VERB, CHUNKING, and private extensions and are not supported. Unsupported commands are translated into Xs, which are rejected by the internal server. This results in a message such as “500 Command unknown: 'XXX'.” Incomplete commands are discarded.

The ESMTP inspection engine changes the characters in the server SMTP banner to asterisks except for the “2”, “0”, “0” characters. Carriage return (CR) and linefeed (LF) characters are ignored.

With SMTP inspection enabled, a Telnet session used for interactive SMTP may hang if the following rules are not observed: SMTP commands must be at least four characters in length; must be terminated with carriage return and line feed; and must wait for a response before issuing the next reply.

An SMTP server responds to client requests with numeric reply codes and optional human-readable strings. SMTP application inspection controls and reduces the commands that the user can use as well as the messages that the server returns. SMTP inspection performs three primary tasks:

- Restricts SMTP requests to seven basic SMTP commands and eight extended commands.
- Monitors the SMTP command-response sequence.
- Generates an audit trail—Audit record 108002 is generated when invalid character embedded in the mail address is replaced. For more information, see RFC 821.

SMTP inspection monitors the command and response sequence for the following anomalous signatures:

- Truncated commands.
- Incorrect command termination (not terminated with <CR><LR>).
- The MAIL and RCPT commands specify who are the sender and the receiver of the mail. Mail addresses are scanned for strange characters. The pipeline character (|) is deleted (changed to a blank space) and "<" ,">" are only allowed if they are used to define a mail address (">" must be preceded by "<").
- Unexpected transition by the SMTP server.
- For unknown commands, the ASA changes all the characters in the packet to X. In this case, the server generates an error code to the client. Because of the change in the packed, the TCP checksum has to be recalculated or adjusted.
- TCP stream editing.
- Command pipelining.

Select ESMTP Map

The Select ESMTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select ESMTP Map

The Select ESMTP Map dialog box lets you select or create a new ESMTP map. An ESMTP map lets you change the configuration values used for ESMTP application inspection. The Select ESMTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default ESMTP inspection map—Specifies to use the default ESMTP map.
- Select an ESMTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

ESMTP Inspect Map

The ESMTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP

The ESMTP pane lets you view previously configured ESMTP application inspection maps. An ESMTP map lets you change the default configuration values used for ESMTP application inspection.

Since ESMTP traffic can be a main source of attack from spam, phising, malformed messages, buffer overflows, and buffer underflows, detailed packet inspection and control of ESMTP traffic are supported. Application security and protocol conformance enforce the sanity of the ESMTP message as well as detect several attacks, block senders and receivers, and block mail relay.

Fields

- ESMTP Inspect Maps—Table that lists the defined ESMTP inspect maps.
- Add—Configures a new ESMTP inspect map. To edit an ESMTP inspect map, choose the ESMTP entry in the ESMTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the ESMTP Inspect Maps table.
- Security Level—Select the security level (high, medium, or low).
  - Low—Default.

    Log if command line length is greater than 512

    Log if command recipient count is greater than 100

    Log if body line length is greater than 1000

    Log if sender address length is greater than 320

    Log if MIME file name length is greater than 255
  - Medium

    Obfuscate Server Banner

    Drop Connections if command line length is greater than 512

    Drop Connections if command recipient count is greater than 100

    Drop Connections if body line length is greater than 1000

    Drop Connections if sender address length is greater than 320

    Drop Connections if MIME file name length is greater than 255
  - High

    Obfuscate Server Banner

    Drop Connections if command line length is greater than 512

    Drop Connections if command recipient count is greater than 100

    Drop Connections if body line length is greater than 1000

    Drop Connections and log if sender address length is greater than 320

    Drop Connections and log if MIME file name length is greater than 255
  - MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.
  - Customize—Opens the Add/Edit ESMTP Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Low.

MIME File Type Filtering

The MIME File Type Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > MIME File Type Filtering

The MIME File Type Filtering dialog box lets you configure the settings for a MIME file type filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add MIME File Type Filter dialog box to add a MIME file type filter.
- Edit—Opens the Edit MIME File Type Filter dialog box to edit a MIME file type filter.
- Delete—Deletes a MIME file type filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit ESMTP Policy Map (Security Level)

The Add/Edit ESMTP Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Basic View

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTPS map is shown.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
  - Low—Default.

    Log if command line length is greater than 512

    Log if command recipient count is greater than 100

    Log if body line length is greater than 1000

    Log if sender address length is greater than 320

    Log if MIME file name length is greater than 255
  - Medium

    Obfuscate Server Banner

    Drop Connections if command line length is greater than 512

    Drop Connections if command recipient count is greater than 100

    Drop Connections if body line length is greater than 1000

    Drop Connections if sender address length is greater than 320

    Drop Connections if MIME file name length is greater than 255
  - High

Obfuscate Server Banner

Drop Connections if command line length is greater than 512

Drop Connections if command recipient count is greater than 100

Drop Connections if body line length is greater than 1000

Drop Connections and log if sender address length is greater than 320

Drop Connections and log if MIME file name length is greater than 255

  - MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.
  - Default Level—Sets the security level back to the default level of Low.
- Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit ESMTP Policy Map (Details)

The Add/Edit ESMTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Advanced View

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTP map is shown.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Shows the security level and mime file type filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the ESMTP inspect map.
  - Mask server banner—Enforces banner obfuscation.
  - Configure Mail Relay—Enables ESMTP mail relay.

    Domain Name—Specifies a local domain.

    Action—Drop connection or log.

    Log—Enable or disable.
- Inspections—Tab that shows you the ESMTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the ESMTP inspection.
  - Value—Shows the value to match in the ESMTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add ESMTP Inspect dialog box to add an ESMTP inspection.
  - Edit—Opens the Edit ESMTP Inspect dialog box to edit an ESMTP inspection.
  - Delete—Deletes an ESMTP inspection.
  - Move Up—Moves an inspection up in the list.

- Move Down—Moves an inspection down in the list.

Add/Edit ESMTP Inspect

The Add/Edit ESMTP Inspect dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Advanced View > Add/Edit ESMTP Inspect

The Add/Edit ESMTP Inspect dialog box lets you define the match criterion and value for the ESMTP inspect map.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.
- Criterion—Specifies which criterion of ESMTP traffic to match.
  - Body Length—Match body length at specified length in bytes.
  - Body Line Length—Match body line length matching at specified length in bytes.
  - Commands—Match commands exchanged in the ESMTP protocol.
  - Command Recipient Count—Match command recipient count greater than number specified.
  - Command Line Length—Match command line length greater than length specified in bytes.
  - EHLO Reply Parameters—Match an ESMTP ehlo reply parameter.
  - Header Length—Match header length at length specified in bytes.
  - Header To Fields Count—Match header To fields count greater than number specified.
  - Invalid Recipients Count—Match invalid recipients count greater than number specified.
  - MIME File Type—Match MIME file type.
  - MIME Filename Length—Match MIME filename.
  - MIME Encoding—Match MIME encoding.
  - Sender Address—Match sender email address.
  - Sender Address Length—Match sender email address length.
- Body Length Criterion Values—Specifies the value details for body length match.
  - Greater Than Length—Body length in bytes.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- Body Line Length Criterion Values—Specifies the value details for body line length match.
  - Greater Than Length—Body line length in bytes.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- Commands Criterion Values—Specifies the value details for command match.
  - Available Commands Table:

AUTH

DATA

EHLO

ETRN

HELO

HELP

MAIL

NOOP

QUIT

RCPT

RSET

SAML

SOML

VRFY

  - Add—Adds the selected command from the Available Commands table to the Selected Commands table.
  - Remove—Removes the selected command from the Selected Commands table.
  - Primary Action—Mask, Reset, Drop Connection, None, Limit Rate (pps).
  - Log—Enable or disable.
  - Rate Limit—Do not limit rate, Limit Rate (pps).
- Command Recipient Count Criterion Values—Specifies the value details for command recipient count match.
  - Greater Than Count—Specify command recipient count.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- Command Line Length Criterion Values—Specifies the value details for command line length.
  - Greater Than Length—Command line length in bytes.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- EHLO Reply Parameters Criterion Values—Specifies the value details for EHLO reply parameters match.
  - Available Parameters Table:

    8bitmime

    auth

    binarymime

    checkpoint

    dsn

    ecode

etrn

others

pipelining

size

vrfy

  - Add—Adds the selected parameter from the Available Parameters table to the Selected Parameters table.
  - Remove—Removes the selected command from the Selected Commands table.
  - Action—Reset, Drop Connection, Mask, Log.
  - Log—Enable or disable.
- Header Length Criterion Values—Specifies the value details for header length match.
  - Greater Than Length—Header length in bytes.
  - Action—Reset, Drop Connection, Mask, Log.
  - Log—Enable or disable.
- Header To Fields Count Criterion Values—Specifies the value details for header To fields count match.
  - Greater Than Count—Specify command recipient count.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- Invalid Recipients Count Criterion Values—Specifies the value details for invalid recipients count match.
  - Greater Than Count—Specify command recipient count.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- MIME File Type Criterion Values—Specifies the value details for MIME file type match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.
- MIME Filename Length Criterion Values—Specifies the value details for MIME filename length match.
  - Greater Than Length—MIME filename length in bytes.
  - Action—Reset, Drop Connection, Log.
  - Log—Enable or disable.
- MIME Encoding Criterion Values—Specifies the value details for MIME encoding match.

  - Available Encodings table

    7bit

    8bit

    base64

    binary

    others

    quoted-printable
  - Add—Adds the selected parameter from the Available Encodings table to the Selected Encodings table.
  - Remove—Removes the selected command from the Selected Commands table.
  - Action—Reset, Drop Connection, Log.
  - Log—Enable or disable.
- Sender Address Criterion Values—Specifies the value details for sender address match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
  - Action—Reset, Drop Connection, Log.
  - Log—Enable or disable.
- Sender Address Length Criterion Values—Specifies the value details for sender address length match.
  - Greater Than Length—Sender address length in bytes.
  - Action—Reset, Drop Connection, Log.
  - Log—Enable or disable.

TFTP Inspection

TFTP inspection is enabled by default.

TFTP, described in RFC 1350, is a simple protocol to read and write files between a TFTP server and client.

The ASA inspects TFTP traffic and dynamically creates connections and translations, if necessary, to permit file transfer between a TFTP client and server. Specifically, the inspection engine inspects TFTP read request (RRQ), write request (WRQ), and error notification (ERROR).

A dynamic secondary channel and a PAT translation, if necessary, are allocated on a reception of a valid read (RRQ) or write (WRQ) request. This secondary channel is subsequently used by TFTP for file transfer or error notification.

Only the TFTP server can initiate traffic over the secondary channel, and at most one incomplete secondary channel can exist between the TFTP client and server. An error notification from the server closes the secondary channel.

TFTP inspection must be enabled if static PAT is used to redirect TFTP traffic.

Configuring Inspection for Voice and Video Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- CTIQBE Inspection, page 59-1
- H.323 Inspection, page 59-2
- MGCP Inspection, page 59-12
- RTSP Inspection, page 59-16
- SIP Inspection, page 59-20
- Skinny (SCCP) Inspection, page 59-31

CTIQBE Inspection

This section describes CTIQBE application inspection. This section includes the following topics:

- CTIQBE Inspection Overview, page 59-1
- Limitations and Restrictions, page 59-2

CTIQBE Inspection Overview

CTIQBE protocol inspection supports NAT, PAT, and bidirectional NAT. This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to work successfully with Cisco CallManager for call setup across the ASA.

TAPI and JTAPI are used by many Cisco VoIP applications. CTIQBE is used by Cisco TSP to communicate with Cisco CallManager.

Limitations and Restrictions

The following summarizes limitations that apply when using CTIQBE application inspection:

- CTIQBE application inspection does not support configurations with the **alias** command.
- Stateful failover of CTIQBE calls is not supported.
- Debugging CTIQBE inspection may delay message transmission, which may have a performance impact in a real-time environment. When you enable this debugging or logging and Cisco IP SoftPhone seems unable to complete call setup through the ASA, increase the timeout values in the Cisco TSP settings on the system running Cisco IP SoftPhone.

The following summarizes special considerations when using CTIQBE application inspection in specific scenarios:

- If two Cisco IP SoftPhones are registered with different Cisco CallManagers, which are connected to different interfaces of the ASA, calls between these two phones fails.
- When Cisco CallManager is located on the higher security interface compared to Cisco IP SoftPhones, if NAT or outside NAT is required for the Cisco CallManager IP address, the mapping must be static as Cisco IP SoftPhone requires the Cisco CallManager IP address to be specified explicitly in its Cisco TSP configuration on the PC.
- When using PAT or Outside PAT, if the Cisco CallManager IP address is to be translated, its TCP port 2748 must be statically mapped to the same port of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not user-configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

H.323 Inspection

This section describes the H.323 application inspection. This section includes the following topics:

- H.323 Inspection Overview, page 59-3
- How H.323 Works, page 59-3
- H.239 Support in H.245 Messages, page 59-4
- Limitations and Restrictions, page 59-4
- Select H.323 Map, page 59-5
- H.323 Class Map, page 59-5
- Add/Edit H.323 Traffic Class Map, page 59-6
- Add/Edit H.323 Match Criterion, page 59-6
- H.323 Inspect Map, page 59-7
- Phone Number Filtering, page 59-8
- Add/Edit H.323 Policy Map (Security Level), page 59-8
- Add/Edit H.323 Policy Map (Details), page 59-9
- Add/Edit HSI Group, page 59-11
- Add/Edit H.323 Map, page 59-11

H.323 Inspection Overview

H.323 inspection provides support for H.323 compliant applications such as Cisco CallManager and VocalTec Gatekeeper. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The ASA supports H.323 through Version 6, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H.323 inspection enabled, the ASA supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the ASA.

The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the ASA uses an ASN.1 decoder to decode the H.323 messages.
- Dynamically allocate the negotiated H.245 and RTP/RTCP connections.

How H.323 Works

The H.323 collection of protocols collectively may use up to two TCP connection and four to eight UDP connections. FastConnect uses only one TCP connection, and RAS uses a single UDP connection for registration, admissions, and status.

An H.323 client can initially establish a TCP connection to an H.323 server using TCP port 1720 to request Q.931 call setup. As part of the call setup process, the H.323 terminal supplies a port number to the client to use for an H.245 TCP connection. In environments where H.323 gatekeeper is in use, the initial packet is transmitted using UDP.

H.323 inspection monitors the Q.931 TCP connection to determine the H.245 port number. If the H.323 terminals are not using FastConnect, the ASA dynamically allocates the H.245 connection based on the inspection of the H.225 messages.

Note The H.225 connection can also be dynamically allocated when using RAS.

Within each H.245 message, the H.323 endpoints exchange port numbers that are used for subsequent UDP data streams. H.323 inspection inspects the H.245 messages to identify these ports and dynamically creates connections for the media exchange. RTP uses the negotiated port number, while RTCP uses the next higher port number.

The H.323 control channel handles H.225 and H.245 and H.323 RAS. H.323 inspection uses the following ports.

- 1718—Gate Keeper Discovery UDP port
- 1719—RAS UDP port
- 1720—TCP Control Port

You must permit traffic for the well-known H.323 port 1719 for RAS signaling. Additionally, you must permit traffic for the well-known H.323 port 1720 for the H.225 call signaling; however, the H.245 signaling ports are negotiated between the endpoints in the H.225 signaling. When an H.323 gatekeeper is used, the ASA opens an H.225 connection based on inspection of the ACF and RCF nmessages.

After inspecting the H.225 messages, the ASA opens the H.245 channel and then inspects traffic sent over the H.245 channel as well. All H.245 messages passing through the ASA undergo H.245 application inspection, which translates embedded IP addresses and opens the media channels negotiated in H.245 messages.

The H.323 ITU standard requires that a TPKT header, defining the length of the message, precede the H.225 and H.245, before being passed on to the reliable connection. Because the TPKT header does not necessarily need to be sent in the same TCP packet as H.225 and H.245 messages, the ASA must remember the TPKT length to process and decode the messages properly. For each connection, the ASA keeps a record that contains the TPKT length for the next expected message.

If the ASA needs to perform NAT on IP addresses in messages, it changes the checksum, the UUIE length, and the TPKT, if it is included in the TCP packet with the H.225 message. If the TPKT is sent in a separate TCP packet, the ASA proxy ACKs that TPKT and appends a new TPKT to the H.245 message with the new length.

Note The ASA does not support TCP options in the Proxy ACK for the TPKT.

Each UDP connection with a packet going through H.323 inspection is marked as an H.323 connection and times out with the H.323 timeout as configured in the Configuration > Firewall > Advanced > Global Timeouts pane.

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The ASA includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled.

H.239 Support in H.245 Messages

The ASA sits between two H.323 endpoints. When the two H.323 endpoints set up a telepresentation session so that the endpoints can send and receive a data presentation, such as spreadsheet data, the ASA ensure successful H.239 negotiation between the endpoints.

H.239 is a standar that provides the ability for H.300 series endpoints to open an additional video channel in a single call. In a call, an endpoint (such as a video phone), sends a channel for video and a channel for data presentation. The H.239 negotiation occurs on the H.245 channel.

The ASA opens pinholes for the additional media channel and the media control channel. The endpoints use open logical channel message (OLC) to signal a new channel creation. The message extension is part of H.245 version 13.

The decoding and encoding of of the telepresentation session is enabled by default. H.239 encoding and decoding is preformed by ASN.1 coder.

Limitations and Restrictions

The following are some of the known issues and limitations when using H.323 application inspection:

- 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 ASA.
- If you configure a network static address where the network static address is the same as a third-party netmask and address, then any outbound H.323 connection fails.
- Configuring both H.323 inspection and communication in and out of the same interface on the ASA is not supported. When you configure both these features, the ASA cannot correctly establish NetMeeting calls because it modifies the H.225 setup message incorrectly by changing the destCallSignalAddress field to point to itself rather than the NetMeeting destination endpoint.

  To workaround this limitation, perform one of the following actions:

  - Configure the ASA so that either H.323 inspection or communication in and out of the same interface, but not both, is set up on the device.
  - Configure inside to inside traffic to use NAT 0.

Select H.323 Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select H.323 Map

The Select H.323 Map dialog box lets you select or create a new H.323 map. An H.323 map lets you change the configuration values used for H.323 application inspection. The Select H.323 Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default H.323 inspection map—Specifies to use the default H.323 map.
- Select an H.323 map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

H.323 Class Map

Configuration > Global Objects > Class Maps > H.323

The H.323 Class Map pane lets you configure H.323 class maps for H.323 inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the H.323 class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the H.323 class map.

  - Value—Shows the value to match in the H.323 class map.
- Description—Shows the description of the class map.
- Add—Adds an H.323 class map.
- Edit—Edits an H.323 class map.
- Delete—Deletes an H.323 class map.

Add/Edit H.323 Traffic Class Map

Configuration > Global Objects > Class Maps > H.323 > Add/Edit H.323 Traffic Class Map

The Add/Edit H.323 Traffic Class Map dialog box lets you define a H.323 class map.

Fields

- Name—Enter the name of the H.323 class map, up to 40 characters in length.
- Description—Enter the description of the H.323 class map.
- Add—Adds an H.323 class map.
- Edit—Edits an H.323 class map.
- Delete—Deletes an H.323 class map.

Add/Edit H.323 Match Criterion

Configuration > Global Objects > Class Maps > H.323 > Add/Edit H.323 Traffic Class Map > Add/Edit H.323 Match Criterion

The Add/Edit H.323 Match Criterion dialog box lets you define the match criterion and value for the H.323 class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of H.323 traffic to match.
  - Called Party—Match the called party.
  - Calling Party—Match the calling party.
  - Media Type—Match the media type.
- Called Party Criterion Values—Specifies to match on the H.323 called party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Calling Party Criterion Values—Specifies to match on the H.323 calling party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Media Type Criterion Values—Specifies which media type to match.
  - Audio—Match audio type.
  - Video—Match video type.
  - Data—Match data type.

H.323 Inspect Map

Configuration > Global Objects > Inspect Maps > H.323

The H.323 pane lets you view previously configured H.323 application inspection maps. An H.323 map lets you change the default configuration values used for H.323 application inspection.

H.323 inspection supports RAS, H.225, and H.245, and its functionality translates all embedded IP addresses and ports. It performs state tracking and filtering and can do a cascade of inspect function activation. H.323 inspection supports phone number filtering, dynamic T.120 control, H.245 tunneling control, HSI groups, protocol state tracking, H.323 call duration enforcement, and audio/video control.

Fields

- H.323 Inspect Maps—Table that lists the defined H.323 inspect maps.
- Add—Configures a new H.323 inspect map. To edit an H.323 inspect map, choose the H.323 entry in the H.323 Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the H.323 Inspect Maps table.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance not enforced
  - Medium

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: no

- High

  State Checking h225 Enabled

  State Checking ras Enabled

  Call Party Number Enabled

  Call duration Limit 1:00:00

  RTP conformance enforced

  Limit payload to audio or video, based on the signaling exchange: yes
- Phone Number Filtering—Opens the Phone Number Filtering dialog box to configure phone number filters.
- Customize—Opens the Add/Edit H.323 Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.

Phone Number Filtering

Configuration > Global Objects > Inspect Maps > H323 > Phone Number Filtering

The Phone Number Filtering dialog box lets you configure the settings for a phone number filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add Phone Number Filter dialog box to add a phone number filter.
- Edit—Opens the Edit Phone Number Filter dialog box to edit a phone number filter.
- Delete—Deletes a phone number filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit H.323 Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Basic View

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.
- Description—Enter the description of the H323 map, up to 200 characters in length.
- Security Level—Select the security level (low, medium, or high).

  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance not enforced
  - Medium

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: no
  - High

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Enabled

    Call duration Limit 1:00:00

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: yes
  - Phone Number Filtering—Opens the Phone Number Filtering dialog box which lets you configure the settings for a phone number filter.
  - Default Level—Sets the security level back to the default.
- Details—Shows the State Checking, Call Attributes, Tunneling and Protocol Conformance, HSI Group Parameters, and Inspections tabs to configure additional settings.

Add/Edit H.323 Policy Map (Details)

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Advanced View

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.
- Description—Enter the description of the H.323 map, up to 200 characters in length.
- Security Level—Shows the security level and phone number filtering settings to configure.
- State Checking—Tab that lets you configure state checking parameters for the H.323 inspect map.
  - Check state transition of H.225 messages—Enforces H.323 state checking on H.225 messages.
  - Check state transition of RAS messages—Enforces H.323 state checking on RAS messages.

  - Check RFC messages and open pinholes for call signal addresses in RFQ messages

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The ASA includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled. You can enable this option by setting the option in the H.323 Inspect Map.

- Call Attributes—Tab that lets you configure call attributes parameters for the H.323 inspect map.
  - Enforce call duration limit—Enforces the absolute limit on a call.

    Call Duration Limit—Time limit for the call (hh:mm:ss).
  - Enforce presence of calling and called party numbers—Enforces sending call party numbers during call setup.
- Tunneling and Protocol Conformance—Tab that lets you configure tunneling and protocol conformance parameters for the H.323 inspect map.
  - Check for H.245 tunneling—Allows H.245 tunneling.

    Action—Drop connection or log.
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio or video based on the signaling exchange.
- HSI Group Parameters—Tab that lets you configure an HSI group.
  - HSI Group ID—Shows the HSI Group ID.
  - IP Address—Shows the HSI Group IP address.
  - Endpoints—Shows the HSI Group endpoints.
  - Add—Opens the Add HSI Group dialog box to add an HSI group.
  - Edit—Opens the Edit HSI Group dialog box to edit an HSI group.
  - Delete—Deletes an HSI group.
- Inspections—Tab that shows you the H.323 inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the H.323 inspection.
  - Value—Shows the value to match in the H.323 inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add H.323 Inspect dialog box to add an H.323 inspection.
  - Edit—Opens the Edit H.323 Inspect dialog box to edit an H.323 inspection.
  - Delete—Deletes an H.323 inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit HSI Group

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Advanced View > Add/Edit HSI Group

The Add/Edit HSI Group dialog box lets you configure HSI Groups.

Fields

- Group ID—Enter the HSI group ID.
- IP Address—Enter the HSI IP address.
- Endpoints—Lets you configure the IP address and interface of the endpoints.
  - IP Address—Enter an endpoint IP address.
  - Interface—Specifies an endpoint interface.
- Add—Adds the HSI group defined.
- Delete—Deletes the selected HSI group.

Add/Edit H.323 Map

Configuration > Global Objects > Inspect Maps > H232 > H323 Inspect Map > Advanced View > Add/Edit H323 Inspect

The Add/Edit H.323 Inspect dialog box lets you define the match criterion and value for the H.323 inspect map.

Fields

- Single Match—Specifies that the H.323 inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of H.323 traffic to match.
  - Called Party—Match the called party.
  - Calling Party—Match the calling party.
  - Media Type—Match the media type.
- Called Party Criterion Values—Specifies to match on the H.323 called party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match on the H.323 calling party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Media Type Criterion Values—Specifies which media type to match.
  - Audio—Match audio type.
  - Video—Match video type.
  - Data—Match data type.
- Multiple Matches—Specifies multiple matches for the H.323 inspection.
  - H323 Traffic Class—Specifies the H.323 traffic class match.
  - Manage—Opens the Manage H323 Class Maps dialog box to add, edit, or delete H.323 Class Maps.
- Action—Drop packet, drop connection, or reset.

MGCP Inspection

This section describes MGCP application inspection. This section includes the following topics:

- MGCP Inspection Overview, page 59-12
- Select MGCP Map, page 59-14
- MGCP Inspect Map, page 59-14
- Gateways and Call Agents, page 59-15
- Add/Edit MGCP Policy Map, page 59-15
- Add/Edit MGCP Group, page 59-16

MGCP Inspection Overview

MGCP is a master/slave protocol used to control media gateways from external call control elements called media gateway controllers or call agents. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Using NAT and PAT with MGCP lets you support a large number of devices on an internal network with a limited set of external (global) addresses. Examples of media gateways are:

- Trunking gateways, that interface between the telephone network and a Voice over IP network. Such gateways typically manage a large number of digital circuits.
- Residential gateways, that provide a traditional analog (RJ11) interface to a Voice over IP network. Examples of residential gateways include cable modem/cable set-top boxes, xDSL devices, broad-band wireless devices.
- Business gateways, that provide a traditional digital PBX interface or an integrated soft PBX interface to a Voice over IP network.

Note To avoid policy failure when upgrading from ASA version 7.1, all layer 7 and layer 3 policies must have distinct names. For instance, a previously configured policy map with the same name as a previously configured MGCP map must be changed before the upgrade.

MGCP messages are transmitted over UDP. A response is sent back to the source address (IP address and UDP port number) of the command, but the response may not arrive from the same address as the command was sent to. This can happen when multiple call agents are being used in a failover configuration and the call agent that received the command has passed control to a backup call agent, which then sends the response. Figure 59-1 illustrates how NAT can be used with MGCP.

Figure 59-1 Using NAT with MGCP

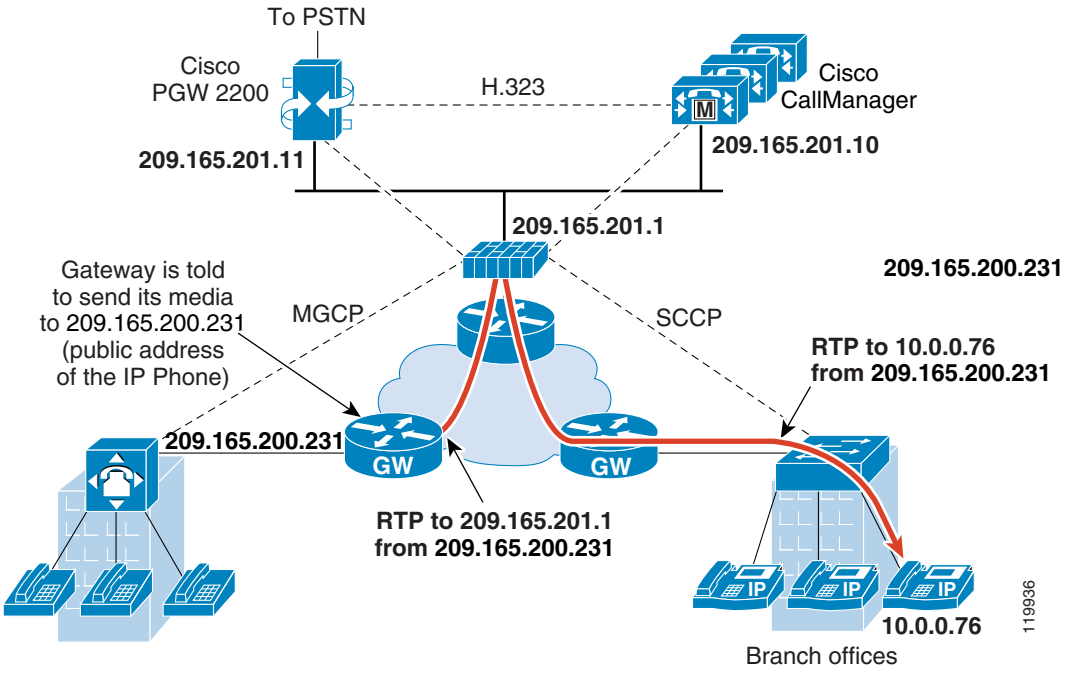

MGCP endpoints are physical or virtual sources and destinations for data. Media gateways contain endpoints on which the call agent can create, modify and delete connections to establish and control media sessions with other multimedia endpoints. Also, the call agent can instruct the endpoints to detect certain events and generate signals. The endpoints automatically communicate changes in service state to the call agent.

MGCP transactions are composed of a command and a mandatory response. There are eight types of commands:

- CreateConnection
- ModifyConnection
- DeleteConnection
- NotificationRequest
- Notify
- AuditEndpoint
- AuditConnection

- RestartInProgress

The first four commands are sent by the call agent to the gateway. The Notify command is sent by the gateway to the call agent. The gateway may also send a DeleteConnection. The registration of the MGCP gateway with the call agent is achieved by the RestartInProgress command. The AuditEndpoint and the AuditConnection commands are sent by the call agent to the gateway.

All commands are composed of a Command header, optionally followed by a session description. All responses are composed of a Response header, optionally followed by a session description.

- The port on which the gateway receives commands from the call agent. Gateways usually listen to UDP port 2427.
- The port on which the call agent receives commands from the gateway. Call agents usually listen to UDP port 2727.

Note MGCP inspection does not support the use of different IP addresses for MGCP signaling and RTP data. A common and recommended practice is to send RTP data from a resilient IP address, such as a loopback or virtual IP address; however, the ASA requires the RTP data to come from the same address as MGCP signalling.

Select MGCP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select MGCP Map

The Select MGCP Map dialog box lets you select or create a new MGCP map. An MGCP map lets you change the configuration values used for MGCP application inspection. The Select MGCP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default MGCP inspection map—Specifies to use the default MGCP map.
- Select an MGCP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

MGCP Inspect Map

Configuration > Global Objects > Inspect Maps > MGCP

The MGCP pane lets you view previously configured MGCP application inspection maps. An MGCP map lets you change the default configuration values used for MGCP application inspection. You can use an MGCP map to manage connections between VoIP devices and MGCP call agents.

Fields

- MGCP Inspect Maps—Table that lists the defined MGCP inspect maps.
- Add—Configures a new MGCP inspect map.
- Edit—Edits the selected MGCP entry in the MGCP Inspect Maps table.
- Delete—Deletes the inspect map selected in the MGCP Inspect Maps table.

Gateways and Call Agents

Configuration > Global Objects > Inspect Maps > MGCP > Gateways and Call Agents

The Gateways and Call Agents dialog box lets you configure groups of gateways and call agents for the map.

Fields

- Group ID—Identifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The gateway IP address can only be associated with one group ID. You cannot use the same gateway with different group IDs. The valid range is from 0 to 2147483647Criterion—Shows the criterion of the inspection.
- Gateways—Identifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.
- Call Agents—Identifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.
- Add—Displays the Add MGCP dialog box, which you can use to define a new application inspection map.
- Edit—Displays the Edit MGCP dialog box, which you can use to modify the application inspection map selected in the application inspection map table.
- Delete—Deletes the application inspection map selected in the application inspection map table.

Add/Edit MGCP Policy Map

Configuration > Global Objects > Inspect Maps > MGCP > MGCP Inspect Map > View

The Add/Edit MGCP Policy Map pane lets you configure the command queue, gateway, and call agent settings for MGCP application inspection maps.

Fields

- Name—When adding an MGCP map, enter the name of the MGCP map. When editing an MGCP map, the name of the previously configured MGCP map is shown.
- Description—Enter the description of the MGCP map, up to 200 characters in length.
- Command Queue—Tab that lets you specify the permitted queue size for MGCP commands.
  - Command Queue Size—Specifies the maximum number of commands to queue. The valid range is from 1 to 2147483647.
- Gateways and Call Agents—Tab that lets you configure groups of gateways and call agents for this map.
  - Group ID—Identifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The gateway IP address can only be associated with one group ID. You cannot use the same gateway with different group IDs. The valid range is from 0 to 2147483647Criterion—Shows the criterion of the inspection.

- Gateways—Identifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.
- Call Agents—Identifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.
- Add—Displays the Add MGCP Group dialog box, which you can use to define a new MGCP group of gateways and call agents.
- Edit—Displays the Edit MGCP dialog box, which you can use to modify the MGCP group selected in the Gateways and Call Agents table.
- Delete—Deletes the MGCP group selected in the Gateways and Call Agents table.

Add/Edit MGCP Group

Configuration > Global Objects > Inspect Maps > MGCP > Add/Edit MGCP Group

The Add/Edit MGCP Group dialog box lets you define the configuration of an MGCP group that will be used when MGCP application inspection is enabled.

Fields

- Group ID—Specifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The valid range is from 0 to 2147483647.
  - Gateway to Be Added—Specifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.
  - Add—Adds the specified IP address to the IP address table.
  - Delete—Deletes the selected IP address from the IP address table.
  - IP Address—Lists the IP addresses of the gateways in the call agent group.
- Call Agents
  - Call Agent to Be Added—Specifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.
  - Add—Adds the specified IP address to the IP address table.
  - Delete—Deletes the selected IP address from the IP address table.
  - IP Address—Lists the IP addresses of the call agents in the call agent group.

RTSP Inspection

This section describes RTSP application inspection. This section includes the following topics:

- RTSP Inspection Overview, page 59-17

- Using RealPlayer, page 59-17
- Restrictions and Limitations, page 59-18
- Select RTSP Map, page 59-18
- RTSP Inspect Map, page 59-18
- Add/Edit RTSP Policy Map, page 59-18
- Add/Edit RTSP Inspect, page 59-19

RTSP Inspection Overview

The RTSP inspection engine lets the ASA pass RTSP packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections.

Note For Cisco IP/TV, use RTSP TCP port 554 and TCP 8554.

RTSP applications use the well-known port 554 with TCP (rarely UDP) as a control channel. The ASA only supports TCP, in conformity with RFC 2326. This TCP control channel is used to negotiate the data channels that is used to transmit audio/video traffic, depending on the transport mode that is configured on the client.

The supported RDT transports are: rtp/avp, rtp/avp/udp, x-real-rdt, x-real-rdt/udp, and x-pn-tng/udp.

The ASA parses Setup response messages with a status code of 200. If the response message is travelling inbound, the server is outside relative to the ASA and dynamic channels need to be opened for connections coming inbound from the server. If the response message is outbound, then the ASA does not need to open dynamic channels.

Because RFC 2326 does not require that the client and server ports must be in the SETUP response message, the ASA keeps state and remembers the client ports in the SETUP message. QuickTime places the client ports in the SETUP message and then the server responds with only the server ports.

RTSP inspection supports PAT or dual-NAT. The ASA provides TCP fragment reassembly support, a scalable parsing routine on RTSP, and security enhancements that protect RTSP traffic.

Using RealPlayer

When using RealPlayer, it is important to properly configure transport mode. For the ASA, add an **access-list** command from the server to the client or vice versa. For RealPlayer, change transport mode by clicking **Options>Preferences>Transport>RTSP Settings**.

If using TCP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use TCP for all content** check boxes. On the ASA, there is no need to configure the inspection engine.

If using UDP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use UDP for static content** check boxes, and for live content not available via Multicast. On the ASA, add an **inspect rtsp** *port* command.

Restrictions and Limitations

The following restrictions apply to the RSTP inspection.

- The ASA does not support multicast RTSP or RTSP messages over UDP.
- The ASA does not have the ability to recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.
- With Cisco IP/TV, the number of translates the ASA performs on the SDP part of the message is proportional to the number of program listings in the Content Manager (each program listing can have at least six embedded IP addresses).
- You can configure NAT for Apple QuickTime 4 or RealPlayer. Cisco IP/TV only works with NAT if the Viewer and Content Manager are on the outside network and the server is on the inside network.

Select RTSP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select NetBIOS Map

The Select RTSP Map dialog box lets you select or create a new RTSP map. An RTSP map lets you change the configuration values used for RTSP application inspection. The Select RTSP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default RTSP inspection map—Specifies to use the default RTSP inspection map.
- Select a RTSP inspect map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

RTSP Inspect Map

Configuration > Global Objects > Inspect Maps > RADIUS

The RTSP pane lets you view previously configured RTSP application inspection maps. An RTSP map lets you change the default configuration values used for RTSP application inspection. You can use an RTSP map to protect RTSP traffic.

Fields

- RTSP Inspect Maps—Table that lists the defined RTSP inspect maps.
- Add—Configures a new RTSP inspect map.
- Edit—Edits the selected RTSP entry in the RTSP Inspect Maps table.
- Delete—Deletes the inspect map selected in the RTSP Inspect Maps table.

Add/Edit RTSP Policy Map

Configuration > Global Objects > Inspect Maps > MGCP > MGCP Inspect Map > View

The Add/Edit RTSP Policy Map pane lets you configure the parameters and inspections settings for RTSP application inspection maps.

Fields

- Name—When adding an RTSP map, enter the name of the RTSP map. When editing an RTSP map, the name of the previously configured RTSP map is shown.
- Description—Enter the description of the RTSP map, up to 200 characters in length.
- Parameters—Tab that lets you restrict usage on reserved ports during media port negotiation, and lets you set the URL length limit.
  - Enforce Reserve Port Protection—Lets you restrict the use of reserved ports during media port negotiation.
  - Maximum URL Length—Specifies the maximum length of the URL allowed in the message. Maximum value is 6000.
- Inspections—Tab that shows you the RTSP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the RTSP inspection.
  - Value—Shows the value to match in the RTSP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add RTSP Inspect dialog box to add a RTSP inspection.
  - Edit—Opens the Edit RTSP Inspect dialog box to edit a RTSP inspection.
  - Delete—Deletes a RTSP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit RTSP Inspect

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View > Add/Edit SIP 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).

SIP Inspection

This section describes SIP application inspection. This section includes the following topics:

- SIP Inspection Overview, page 59-20
- SIP Instant Messaging, page 59-21
- Select SIP Map, page 59-22
- SIP Class Map, page 59-23
- Add/Edit SIP Traffic Class Map, page 59-23
- Add/Edit SIP Match Criterion, page 59-24
- SIP Inspect Map, page 59-26
- Add/Edit SIP Policy Map (Security Level), page 59-27
- Add/Edit SIP Policy Map (Details), page 59-28
- Add/Edit SIP Inspect, page 59-29

SIP Inspection Overview

SIP, as defined by the IETF, enables call handling sessions, particularly two-party audio conferences, or “calls.” SIP works with SDP for call signalling. SDP specifies the ports for the media stream. Using SIP, the ASA can support any SIP VoIP gateways and VoIP proxy servers. SIP and SDP are defined in the following RFCs:

- SIP: Session Initiation Protocol, RFC 3261
- SDP: Session Description Protocol, RFC 2327

To support SIP calls through the ASA, signaling messages for the media connection addresses, media ports, and embryonic connections for the media must be inspected, because while the signaling is sent over a well-known destination port (UDP/TCP 5060), the media streams are dynamically allocated. Also, SIP embeds IP addresses in the user-data portion of the IP packet. SIP inspection applies NAT for these embedded IP addresses.

The following limitations and restrictions apply when using PAT with SIP:

- If a remote endpoint tries to register with a SIP proxy on a network protected by the ASA, the registration fails under very specific conditions, as follows:
  - PAT is configured for the remote endpoint.
  - The SIP registrar server is on the outside network.
  - The port is missing in the contact field in the REGISTER message sent by the endpoint to the proxy server.
- 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 ASA opens RTP/RTCP connections between the two endpoints using these media addresses/ports.

The well-known port 5060 must be used on the initial call setup (INVITE) message; however, subsequent messages may not have this port number. The SIP inspection engine opens signaling connection pinholes, and marks these connections as SIP connections. This is done for the messages to reach the SIP application and be translated.

As a call is set up, the SIP session is in the "transient" state until the media address and media port is received from the called endpoint in a Response message indicating the RTP port the called endpoint listens on. If there is a failure to receive the response messages within one minute, the signaling connection is torn down.

Once the final handshake is made, the call state is moved to active and the signaling connection remains until a BYE message is received.

If an inside endpoint initiates a call to an outside endpoint, a media hole is opened to the outside interface to allow RTP/RTCP UDP packets to flow to the inside endpoint media address and media port specified in the INVITE message from the inside endpoint. Unsolicited RTP/RTCP UDP packets to an inside interface does not traverse the ASA, unless the ASA configuration specifically allows it.

Select SIP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select SIP Map

The Select SIP Map dialog box lets you select or create a new SIP map. A SIP map lets you change the configuration values used for SIP application inspection. The Select SIP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SIP inspection map—Specifies to use the default SIP map.
- Select a SIP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.
- Enable encrypted traffic inspection check box—Select to enable the radio buttons to select a proxy type.
- Proxy Type
  - TLS Proxy radio button—Use TLS Proxy to enable inspection of encrypted traffic.
  - Phone Proxy radio button—Specifies to associate the Phone Proxy with the TLS Proxy that you select from the TLS Proxy Name field.

    Configure button—Opens the Configure the Phone Proxy dialog box so that you can specify or edit Phone Proxy configuration settings.
  - UC-IME Proxy ratio button—Specifies to associate the UC-IME Proxy (Cisco Intercompany Media Engine proxy) with the TLS Proxy that you select from the TLS Proxy Name field.

    Configure button—Opens the Configure the UC-IME Proxy dialog box so that you can specify or edit UC-IME Proxy configuration settings.
- TLS Proxy Name:—Name of existing TLS Proxy.
- Manage—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

Only one TLS proxy can be assigned to the Phone Proxy or UC-IME Proxy at a time. If you configure more than one service policy rule for Phone Proxy or UC-IME Proxy inspection and attempt to assign a different TLS proxy to them, ASDM displays a warning that all other service policy rules with Phone Proxy or UC-IME inspection will be changed to use the latest selected TLS proxy.

The UC-IME Proxy configuration requires two TLS proxies – one for outbound traffic and one for inbound. Rather than associating the TLS proxies directly with the UC-IME Proxy, as is the case with phone proxy, the TLS proxies are associated with it indirectly via SIP inspection rules.

You associate a TLS proxy with the Phone Proxy while defining a SIP inspection action . ASDM will convert the association to the existing phone proxy.

SIP Class Map

Configuration > Global Objects > Class Maps > SIP

The SIP Class Map pane lets you configure SIP class maps for SIP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the SIP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the SIP class map.
  - Value—Shows the value to match in the SIP class map.
- Description—Shows the description of the class map.
- Add—Adds a SIP class map.
- Edit—Edits a SIP class map.
- Delete—Deletes a SIP class map.

Add/Edit SIP Traffic Class Map

Configuration > Global Objects > Class Maps > SIP > Add/Edit SIP Traffic Class Map

The Add/Edit SIP Traffic Class Map dialog box lets you define a SIP class map.

Fields

- Name—Enter the name of the SIP class map, up to 40 characters in length.
- Description—Enter the description of the SIP class map.
- Add—Adds a SIP class map.
- Edit—Edits a SIP class map.
- Delete—Deletes a SIP class map.

Add/Edit SIP Match Criterion

Configuration > Global Objects > Class Maps > SIP > Add/Edit SIP Traffic Class Map > Add/Edit SIP Match Criterion

The Add/Edit SIP Match Criterion dialog box lets you define the match criterion and value for the SIP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of SIP traffic to match.
  - Called Party—Match the called party as specified in the To header.
  - Calling Party—Match the calling party as specified in the From header.
  - Content Length—Match the Content Length header, between 0 and 65536.
  - Content Type—Match the Content Type header.
  - IM Subscriber—Match the SIP IM subscriber.
  - Message Path—Match the SIP Via header.
  - Request Method—Match the SIP request method.
  - Third-Party Registration—Match the requester of a third-party registration.
  - URI Length—Match a URI in the SIP headers, between 0 and 65536.
- Called Party Criterion Values—Specifies to match the called party. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match the calling party. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Content Length Criterion Values—Specifies to match a SIP content header of a length greater than specified.
  - Greater Than Length—Enter a header length value in bytes.
- Content Type Criterion Values—Specifies to match a SIP content header type.

  - SDP—Match an SDP SIP content header type.
  - Regular Expression—Match a regular expression.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- IM Subscriber Criterion Values—Specifies to match the IM subscriber. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Message Path Criterion Values—Specifies to match a SIP Via header. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Request Method Criterion Values—Specifies to match a SIP request method.
  - Request Method—Specifies a request method: ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update.
- Third-Party Registration Criterion Values—Specifies to match the requester of a third-party registration. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- URI Length Criterion Values—Specifies to match a URI of a selected type and greater than the specified length in the SIP headers.
  - URI type—Specifies to match either SIP URI or TEL URI.
  - Greater Than Length—Length in bytes.

SIP Inspect Map

Configuration > Global Objects > Inspect Maps > SIP

The SIP pane lets you view previously configured SIP application inspection maps. A SIP map lets you change the default configuration values used for SIP application inspection.

SIP is a widely used protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. Partially because of its text-based nature and partially because of its flexibility, SIP networks are subject to a large number of security threats.

SIP application inspection provides address translation in message header and body, dynamic opening of ports and basic sanity checks. It also supports application security and protocol conformance, which enforce the sanity of the SIP messages, as well as detect SIP-based attacks.

Fields

- SIP Inspect Maps—Table that lists the defined SIP inspect maps.
- Add—Configures a new SIP inspect map. To edit a SIP inspect map, choose the SIP entry in the SIP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the SIP Inspect Maps table.
- Security Level—Select the security level (high or low).
  - Low—Default.

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Not enforced.

    SIP conformance: Do not perform state checking and header validation.
  - Medium

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No

    SIP conformance: Drop packets that fail state checking.
  - High

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Denied.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Enabled.

Ensure that the number of hops to destination is greater than 0: Enabled.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: Yes

SIP conformance: Drop packets that fail state checking and packets that fail header validation.

- Customize—Opens the Add/Edit SIP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

Add/Edit SIP Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Basic View

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

- Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.
- Description—Enter the description of the SIP map, up to 200 characters in length.
- Security Level—Select the security level (high or low).
  - Low—Default.

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Not enforced.

    SIP conformance: Do not perform state checking and header validation.
  - Medium

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No

    SIP conformance: Drop packets that fail state checking.
  - High

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Denied.

    Hide server's and endpoint's IP addresses: Disabled.

Mask software version and non-SIP URIs: Enabled.

Ensure that the number of hops to destination is greater than 0: Enabled.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: Yes

SIP conformance: Drop packets that fail state checking and packets that fail header validation.

  - Default Level—Sets the security level back to the default.
- Details—Shows additional filtering, IP address privacy, hop count, RTP conformance, SIP conformance, field masking, and inspections settings to configure.

Add/Edit SIP Policy Map (Details)

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

- Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.
- Description—Enter the description of the SIP map, up to 200 characters in length.
- Security Level—Shows the security level settings to configure
- Filtering—Tab that lets you configure the filtering settings for SIP.
  - Enable SIP instant messaging (IM) extensions—Enables Instant Messaging extensions. Default is enabled.
  - Permit non-SIP traffic on SIP port—Permits non-SIP traffic on SIP port. Permitted by default.
- IP Address Privacy—Tab that lets you configure the IP address privacy settings for SIP.
  - Hide server's and endpoint's IP addresses—Enables IP address privacy. Disabled by default.
- Hop Count—Tab that lets you configure the hop count settings for SIP.
  - Ensure that number of hops to destination is greater than 0—Enables check for the value of Max-Forwards header is zero.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.
- RTP Conformance—Tab that lets you configure the RTP conformance settings for SIP.
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces payload type to be audio/video based on the signaling exchange.
- SIP Conformance—Tab that lets you configure the SIP conformance settings for SIP.
  - Enable state transition checking—Enables SIP state checking.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.
  - Enable strict validation of header fields—Enables validation of SIP header fields.

Action—Drop packet, Drop Connection, Reset, Log.

Log—Enable or Disable.

- Field Masking—Tab that lets you configure the field masking settings for SIP.
  - Inspect non-SIP URIs—Enables non-SIP URI inspection in Alert-Info and Call-Info headers.

    Action—Mask or Log.

    Log—Enable or Disable.
  - Inspect server's and endpoint's software version—Inspects SIP endpoint software version in User-Agent and Server headers.

    Action—Mask or Log.

    Log—Enable or Disable.
- Inspections—Tab that shows you the SIP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the SIP inspection.
  - Value—Shows the value to match in the SIP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add SIP Inspect dialog box to add a SIP inspection.
  - Edit—Opens the Edit SIP Inspect dialog box to edit a SIP inspection.
  - Delete—Deletes a SIP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit SIP Inspect

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View > Add/Edit SIP Inspect

The Add/Edit SIP Inspect dialog box lets you define the match criterion and value for the SIP inspect map.

Fields

- Single Match—Specifies that the SIP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.
- Criterion—Specifies which criterion of SIP traffic to match.
  - Called Party—Match a called party as specified in the To header.
  - Calling Party—Match a calling party as specified in the From header.
  - Content Length—Match a content length header.
  - Content Type—Match a content type header.
  - IM Subscriber—Match a SIP IM subscriber.

  - Message Path—Match a SIP Via header.
  - Request Method—Match a SIP request method.
  - Third-Party Registration—Match the requester of a third-party registration.
  - URI Length—Match a URI in the SIP headers.
- Called Party Criterion Values—Specifies to match the called party. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match the calling party. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Content Length Criterion Values—Specifies to match a SIP content header of a length greater than specified.
  - Greater Than Length—Enter a header length value in bytes.
- Content Type Criterion Values—Specifies to match a SIP content header type.
  - SDP—Match an SDP SIP content header type.
  - Regular Expression—Match a regular expression.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- IM Subscriber Criterion Values—Specifies to match the IM subscriber. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Message Path Criterion Values—Specifies to match a SIP Via header. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Request Method Criterion Values—Specifies to match a SIP request method.
  - Request Method—Specifies a request method: ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update.
- Third-Party Registration Criterion Values—Specifies to match the requester of a third-party registration. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- URI Length Criterion Values—Specifies to match a URI in the SIP headers greater than specified length.
  - URI type—Specifies to match either SIP URI or TEL URI.
  - Greater Than Length—Length in bytes.
- Multiple Matches—Specifies multiple matches for the SIP inspection.
  - SIP Traffic Class—Specifies the SIP traffic class match.
  - Manage—Opens the Manage SIP Class Maps dialog box to add, edit, or delete SIP Class Maps.
- Actions—Primary action and log settings.
  - Action—Drop packet, drop connection, reset, log. Note: Limit rate (pps) action is available for request methods invite and register.
  - Log—Enable or disable.

Skinny (SCCP) Inspection

This section describes SCCP application inspection. This section includes the following topics:

- SCCP Inspection Overview, page 59-32
- Supporting Cisco IP Phones, page 59-32
- Restrictions and Limitations, page 59-33
- Select SCCP (Skinny) Map, page 59-33
- SCCP (Skinny) Inspect Map, page 59-34
- Message ID Filtering, page 59-35

- Add/Edit SCCP (Skinny) Policy Map (Security Level), page 59-35
- Add/Edit SCCP (Skinny) Policy Map (Details), page 59-36
- Add/Edit Message ID Filter, page 59-37

SCCP Inspection Overview

Note For specific information about setting up the Phone Proxy on the ASA, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see Chapter 65, "Configuring the Cisco Phone Proxy.".

Skinny (SCCP) is a simplified protocol used in VoIP networks. Cisco IP Phones using SCCP can coexist in an H.323 environment. When used with Cisco CallManager, the SCCP client can interoperate with H.323 compliant terminals.

The ASA supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signalling and media packets can traverse the ASA.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The ASA also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

Note The ASA supports inspection of traffic from Cisco IP Phones running SCCP protocol version 19 and earlier.

Supporting Cisco IP Phones

Note For specific information about setting up the Phone Proxy on the ASA, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see Chapter 65, "Configuring the Cisco Phone Proxy."

In topologies where Cisco CallManager is located on the higher security interface with respect to the Cisco IP Phones, if NAT is required for the Cisco CallManager IP address, the mapping must be **static** as a Cisco IP Phone requires the Cisco CallManager IP address to be specified explicitly in its configuration. An static identity entry allows the Cisco CallManager on the higher security interface to accept registrations from the Cisco IP Phones.

Cisco IP Phones require access to a TFTP server to download the configuration information they need to connect to the Cisco CallManager server.

When the Cisco IP Phones are on a lower security interface compared to the TFTP server, you must use an 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

The following are some of the known issues and limitations when using SCCP application inspection:

- PAT does not work with configurations containing the **alias** command.
- Outside NAT or PAT is *not* supported.

  If the address of an internal Cisco CallManager is configured for NAT or PAT to a different IP address or port, registrations for external Cisco IP Phones fail because the ASA currently does not support NAT or PAT for the file content transferred over TFTP.

  Although the ASA supports NAT of TFTP messages and opens a pinhole for the TFTP file, the ASA cannot translate the Cisco CallManager IP address and port embedded in the Cisco IP Phone configuration files that are transferred by TFTP during phone registration.
- The ASA supports stateful failover of SCCP calls except for calls that are in the middle of call setup.
- When the ASA is running in transparent firewall mode, it blocks SCCP voice traffic because pinholes are not opened for the connection. Displaying debugging messages for SCCP inspection indicate that RTP and RTCP channels are not open. This limication affects communication between IP Phones and the Unified Communications Manager.

  To workaround this limitation, perform one of the following actions to open secondary connections:

  – (Preferred) Install static route entries for IP Phones or convert to routed mode.

  – Configure static ARP entries for IP Phones

    Example:

    ```
    UC Manager---IP Phone 1---Inside ASA (Transparent)---Outside ASA---IP Phone 2
    ```

    In this example, configure a static ARP entry for IP Phone 1 on the ASA to send RTP and RTCP traffic from IP Phone 2 to IP Phone 1.

Select SCCP (Skinny) Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select SCCP Map

The Select SCCP (Skinny) Map dialog box lets you select or create a new SCCP (Skinny) map. An SCCP (Skinny) map lets you change the configuration values used for SCCP (Skinny) application inspection. The Select SCCP (Skinny) Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SCCP (Skinny) inspection map—Specifies to use the default SCCP (Skinny) map.
- Select an SCCP (Skinny) map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.
- Encrypted Traffic Inspection—Lets you specify TLS proxy settings for the inspect map.

- Do not inspect Encrypted Traffic—Disables the inspection of Skinny application inspection.
- Use Phone Proxy to enable inspection of encrypted traffic—Uses the Phone Proxy configured on the ASA to inspect Skinny application traffic. See Chapter 65, "Configuring the Cisco Phone Proxy."
- Use TLS Proxy to enable inspection of encrypted traffic—Specifies to use Transaction Layer Security Proxy to enable inspection of encryped traffic.

  TLS Proxy Name:—Name of existing TLS Proxy.

  New—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

SCCP (Skinny) Inspect Map

Configuration > Global Objects > Inspect Maps > SCCP (Skinny)

The SCCP (Skinny) pane lets you view previously configured SCCP (Skinny) application inspection maps. An SCCP (Skinny) map lets you change the default configuration values used for SCCP (Skinny) application inspection.

Skinny application inspection performs translation of embedded IP address and port numbers within the packet data, and dynamic opening of pinholes. It also performs additional protocol conformance checks and basic state tracking.

Fields

- SCCP (Skinny) Inspect Maps—Table that lists the defined SCCP (Skinny) inspect maps.
- Add—Configures a new SCCP (Skinny) inspect map. To edit an SCCP (Skinny) inspect map, choose the SCCP (Skinny) entry in the SCCP (Skinny) Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the SCCP (Skinny) Inspect Maps table.
- Security Level—Select the security level (high or low).
  - Low—Default.

    Registration: Not enforced.

    Maximum message ID: 0x181.

    Minimum prefix length: 4

    Media timeout: 00:05:00

    Signaling timeout: 01:00:00.

    RTP conformance: Not enforced.
  - Medium

    Registration: Not enforced.

    Maximum message ID: 0x141.

    Minimum prefix length: 4.

    Media timeout: 00:01:00.

    Signaling timeout: 00:05:00.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No.
  - High

Registration: Enforced.

Maximum message ID: 0x141.

Minimum prefix length: 4.

Maximum prefix length: 65536.

Media timeout: 00:01:00.

Signaling timeout: 00:05:00.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: Yes.

- Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
- Customize—Opens the Add/Edit SCCP (Skinny) Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

Message ID Filtering

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > Message ID Filtering

The Message ID Filtering dialog box lets you configure the settings for a message ID filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add Message ID Filtering dialog box to add a message ID filter.
- Edit—Opens the Edit Message ID Filtering dialog box to edit a message ID filter.
- Delete—Deletes a message ID filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit SCCP (Skinny) Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Basic View

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.
- Description—Enter the description of the SCCP (Skinny) map, up to 200 characters in length.

- Security Level—Select the security level (high or low).
  - Low—Default.

    Registration: Not enforced.

    Maximum message ID: 0x181.

    Minimum prefix length: 4

    Media timeout: 00:05:00

    Signaling timeout: 01:00:00.

    RTP conformance: Not enforced.
  - Medium

    Registration: Not enforced.

    Maximum message ID: 0x141.

    Minimum prefix length: 4.

    Media timeout: 00:01:00.

    Signaling timeout: 00:05:00.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No.
  - High

    Registration: Enforced.

    Maximum message ID: 0x141.

    Minimum prefix length: 4.

    Maximum prefix length: 65536.

    Media timeout: 00:01:00.

    Signaling timeout: 00:05:00.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: Yes.
  - Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
  - Default Level—Sets the security level back to the default.
- Details—Shows additional parameter, RTP conformance, and message ID filtering settings to configure.

Add/Edit SCCP (Skinny) Policy Map (Details)

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Advanced View

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the security level and message ID filtering settings to configure.
- Parameters—Tab that lets you configure the parameter settings for SCCP (Skinny).
  - Enforce endpoint registration—Enforce that Skinny endpoints are registered before placing or receiving calls.

    Maximum Message ID—Specify value of maximum SCCP message ID allowed.
  - SCCP Prefix Length—Specifies prefix length value in Skinny messages.

    Minimum Prefix Length—Specify minimum value of SCCP prefix length allowed.

    Maximum Prefix Length—Specify maximum value of SCCP prefix length allowed.
  - Media Timeout—Specify timeout value for media connections.
  - Signaling Timeout—Specify timeout value for signaling connections.
- RTP Conformance—Tab that lets you configure the RTP conformance settings for SCCP (Skinny).
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio/video based on the signaling exchange.
- Message ID Filtering—Tab that lets you configure the message ID filtering settings for SCCP (Skinny).
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the inspection.
  - Value—Shows the value to match in the inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add Message ID Filtering dialog box to add a message ID filter.
  - Edit—Opens the Edit Message ID Filtering dialog box to edit a message ID filter.
  - Delete—Deletes a message ID filter.
  - Move Up—Moves an entry up in the list.
  - Move Down—Moves an entry down in the list.

Add/Edit Message ID Filter

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Advanced View > Add/Edit Message ID Filter

The Add Message ID Filter dialog box lets you configure message ID filters.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.

- Criterion—Specifies which criterion of SCCP (Skinny) traffic to match.
  - Message ID—Match specified message ID.

    Message ID—Specify value of maximum SCCP message ID allowed.
  - Message ID Range—Match specified message ID range.

    Lower Message ID—Specify lower value of SCCP message ID allowed.

    Upper Message ID—Specify upper value of SCCP message ID allowed.
- Action—Drop packet.
- Log—Enable or disable.

Configuring Inspection of Database and Directory Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- ILS Inspection, page 60-1
- SQL*Net Inspection, page 60-2
- Sun RPC Inspection, page 60-3

ILS Inspection

The ILS inspection engine provides NAT support for Microsoft NetMeeting, SiteServer, and Active Directory products that use LDAP to exchange directory information with an ILS server.

The ASA supports NAT for ILS, which is used to register and locate endpoints in the ILS or SiteServer Directory. PAT cannot be supported because only IP addresses are stored by an LDAP database.

For search responses, when the LDAP server is located outside, NAT should be considered to allow internal peers to communicate locally while registered to external LDAP servers. For such search responses, xlates are searched first, and then DNAT entries to obtain the correct address. If both of these searches fail, then the address is not changed. For sites using NAT 0 (no NAT) and not expecting DNAT interaction, we recommend that the inspection engine be turned off to provide better performance.

Additional configuration may be necessary when the ILS server is located inside the ASA border. This would require a hole for outside clients to access the LDAP server on the specified port, typically TCP 389.

Because ILS traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the TCP inactivity interval. By default, this interval is 60 minutes and can be adjusted using the **timeout** command.

ILS/LDAP follows a client/server model with sessions handled over a single TCP connection. Depending on the client's actions, several of these sessions may be created.

During connection negotiation time, a BIND PDU is sent from the client to the server. Once a successful BIND RESPONSE from the server is received, other operational messages may be exchanged (such as ADD, DEL, SEARCH, or MODIFY) to perform operations on the ILS Directory. The ADD REQUEST and SEARCH RESPONSE PDUs may contain IP addresses of NetMeeting peers, used by H.323 (SETUP and CONNECT messages) to establish the NetMeeting sessions. Microsoft NetMeeting v2.X and v3.X provides ILS support.

The ILS inspection performs the following operations:

- Decodes the LDAP REQUEST/RESPONSE PDUs using the BER decode functions
- Parses the LDAP packet
- Extracts IP addresses
- Translates IP addresses as necessary
- Encodes the PDU with translated addresses using BER encode functions
- Copies the newly encoded PDU back to the TCP packet
- Performs incremental TCP checksum and sequence number adjustment

ILS inspection has the following limitations:

- Referral requests and responses are not supported
- Users in multiple directories are not unified
- Single users having multiple identities in multiple directories cannot be recognized by NAT

Note Because H.225 call signalling traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the interval specified by the TCP option in the Configuration > Firewall > Advanced > Global Timeouts pane. By default, this interval is set at 60 minutes.

SQL*Net Inspection

SQL*Net inspection is enabled by default.

The SQL*Net protocol consists of different packet types that the ASA handles to make the data stream appear consistent to the Oracle applications on either side of the ASA.

The default port assignment for SQL*Net is 1521. This is the value used by Oracle for SQL*Net, but this value does not agree with IANA port assignments for Structured Query Language (SQL).

Note Disable SQL*Net inspection when SQL data transfer occurs on the same port as the SQL control TCP port 1521. The security appliance acts as a proxy when SQL*Net inspection is enabled and reduces the client window size from 65000 to about 16000 causing data transfer issues.

The ASA translates all addresses and looks in the packets for all embedded ports to open for SQL*Net Version 1.

For SQL*Net Version 2, all DATA or REDIRECT packets that immediately follow REDIRECT packets with a zero data length will be fixed up.

The packets that need fix-up contain embedded host/port addresses in the following format:

```
(ADDRESS=(PROTOCOL=tcp)(DEV=6)(HOST=a.b.c.d)(PORT=a))
```

SQL*Net Version 2 TNSFrame types (Connect, Accept, Refuse, Resend, and Marker) will not be scanned for addresses to NAT nor will inspection open dynamic connections for any embedded ports in the packet.

SQL*Net Version 2 TNSFrames, Redirect, and Data packets will be scanned for ports to open and addresses to NAT, if preceded by a REDIRECT TNSFrame type with a zero data length for the payload. When the Redirect message with data length zero passes through the ASA, a flag will be set in the connection data structure to expect the Data or Redirect message that follows to be translated and ports to be dynamically opened. If one of the TNS frames in the preceding paragraph arrive after the Redirect message, the flag will be reset.

The SQL*Net inspection engine will recalculate the checksum, change IP, TCP lengths, and readjust Sequence Numbers and Acknowledgment Numbers using the delta of the length of the new and old message.

SQL*Net Version 1 is assumed for all other cases. TNSFrame types (Connect, Accept, Refuse, Resend, Marker, Redirect, and Data) and all packets will be scanned for ports and addresses. Addresses will be translated and port connections will be opened.

Sun RPC Inspection

This section describes Sun RPC application inspection. This section includes the following topics:

- Sun RPC Inspection Overview, page 60-3
- "SUNRPC Server" section on page 60-4
- "Add/Edit SUNRPC Service" section on page 60-4

Sun RPC Inspection Overview

The Sun RPC inspection engine enables or disables application inspection for the Sun RPC protocol. Sun RPC is used by NFS and NIS. Sun RPC services can run on any port. When a client attempts to access an Sun RPC service on a server, it must learn the port that service is running on. It does this by querying the port mapper process, usually rpcbind, on the well-known port of 111.

The client sends the Sun RPC program number of the service and the port mapper process responds with the port number of the service. The client sends its Sun RPC queries to the server, specifying the port identified by the port mapper process. When the server replies, the ASA intercepts this packet and opens both embryonic TCP and UDP connections on that port.

The following limitations apply to Sun RPC inspection:

- NAT or PAT of Sun RPC payload information is not supported.
- Sun RPC inspection supports inbound access lists only. Sun RPC inspection does not support outbound access lists because the inspection engine uses dynamic access lists instead of secondary connections. Dynamic access lists are always added on the ingress direction and not on egress; therefore, this inspection engine does not support outbound access lists. To view the dynamic access lists configured for the ASA, use the **show asp table classify domain permit** command. For information about the **show asp table classify domain permit** command, see the CLI configuration guide.

SUNRPC Server

Configuration > Properties > SUNRPC Server

The Configuration > Firewall > Advanced > SUNRPC Server pane shows which SunRPC services can traverse the ASA and their specific timeout, on a per server basis.

Fields

- Interface—Displays the interface on which the SunRPC server resides.
- IP address—Displays the IP address of the SunRPC server.
- Mask—Displays the subnet mask of the IP Address of the SunRPC server.
- Service ID—Displays the SunRPC program number, or service ID, allowed to traverse the ASA.
- Protocol—Displays the SunRPC transport protocol (TCP or UDP).
- Port—Displays the SunRPC protocol port range.
- Timeout—Displays the idle time after which the access for the SunRPC service traffic is closed.

Add/Edit SUNRPC Service

Configuration > Properties > SUNRPC Server > Add/Edit SUNRPC Service

The Configuration > Firewall > Advanced > SUNRPC Server > Add/Edit SUNRPC Service dialog box lets you specify what SunRPC services are allowed to traverse the ASA and their specific timeout, on a per-server basis.

Fields

- Interface Name—Specifies the interface on which the SunRPC server resides.
- Protocol—Specifies the SunRPC transport protocol (TCP or UDP).
- IP address—Specifies the IP address of the SunRPC server.
- Port—Specifies the SunRPC protocol port range.
- Mask—Specifies the subnet mask of the IP Address of the SunRPC server.
- Timeout—Specifies the idle time after which the access for the SunRPC service traffic is closed. Format is HH:MM:SS.
- Service ID—Specifies the SunRPC program number, or service ID, allowed to traverse the ASA.

Configuring Inspection for Management Application Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- DCERPC Inspection, page 61-1
- GTP Inspection, page 61-4
- RADIUS Accounting Inspection, page 61-10
- RSH Inspection, page 61-13
- SNMP Inspection, page 61-13
- XDMCP Inspection, page 61-15

DCERPC Inspection

This section describes the DCERPC inspection engine. This section includes the following topics:

- DCERPC Overview, page 61-1
- “Select DCERPC Map” section on page 61-2
- “DCERPC Inspect Map” section on page 61-2
- “Add/Edit DCERPC Policy Map” section on page 61-3

DCERPC Overview

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Since a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Note DCERPC inspection only supports communication between the EPM and clients to open pinholes through theASA. Clients using RPC communication that does not use the EPM is not supported with DCERPC inspection.

Select DCERPC Map

Add/Edit Service Policy Rule Wizard > Rule Actions >
Protocol Inspection Tab >Select DCERPC Map

The Select DCERPC Map dialog box lets you select or create a new DCERPC map. A DCERPC map lets you change the configuration values used for DCERPC application inspection. The Select DCERPC Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default DCERPC inspection map—Specifies to use the default DCERPC map.
- Select a DCERPC map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

DCERPC Inspect Map

Configuration > Global Objects > Inspect Maps > DCERPC

The DCERPC pane lets you view previously configured DCERPC application inspection maps. A DCERPC map lets you change the default configuration values used for DCERPC application inspection.

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper (EPM) listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Because a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Fields

- DCERPC Inspect Maps—Table that lists the defined DCERPC inspect maps.
- Add—Configures a new DCERPC inspect map. To edit a DCERPC inspect map, choose the DCERPC entry in the DCERPC Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the DCERPC Inspect Maps table.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: enabled

    Endpoint mapper service lookup timeout: 00:05:00
  - Medium—Default.

    Pinhole timeout: 00:01:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: disabled.
  - High

    Pinhole timeout: 00:01:00

    Endpoint mapper service: enforced

    Endpoint mapper service lookup: disabled
  - Customize—Opens the Add/Edit DCERPC Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.

Add/Edit DCERPC Policy Map

Configuration > Global Objects > Inspect Maps > DCERPC > DCERPC Inspect Map > Basic/Advanced View

The Add/Edit DCERPC Policy Map pane lets you configure the security level and parameters for DCERPC application inspection maps.

Fields

- Name—When adding a DCERPC map, enter the name of the DCERPC map. When editing a DCERPC map, the name of the previously configured DCERPC map is shown.
- Description—Enter the description of the DCERPC map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

Endpoint mapper service: not enforced

Endpoint mapper service lookup: enabled

Endpoint mapper service lookup timeout: 00:05:00

- Medium—Default.

  Pinhole timeout: 00:01:00

  Endpoint mapper service: not enforced

  Endpoint mapper service lookup: disabled.

- High

  Pinhole timeout: 00:01:00

  Endpoint mapper service: enforced

  Endpoint mapper service lookup: disabled

- Default Level—Sets the security level back to the default level of Medium.

- Details—Shows the Parameters to configure additional settings.
  - Pinhole Timeout—Sets the pinhole timeout. Because a client may use the server information returned by the endpoint mapper for multiple connections, the timeout value is configurable based on the client application environment. Range is from 0:0:1 to 1193:0:0. Default is 2 minutes.
  - Enforce endpoint-mapper service—Enforces endpoint mapper service during binding.
  - Enable endpoint-mapper service lookup—Enables the lookup operation of the endpoint mapper service. If disabled, the pinhole timeout is used.

    Enforce Service Lookup Timeout—Enforces the service lookup timeout specified.

    Service Lookup Timeout—Sets the timeout for pinholes from lookup operation.

GTP Inspection

This section describes the GTP inspection engine. This section includes the following topics:

- GTP Inspection Overview, page 61-5
- "Select GTP Map" section on page 61-5
- "GTP Inspect Map" section on page 61-6
- "IMSI Prefix Filtering" section on page 61-7
- "Add/Edit GTP Policy Map (Security Level)" section on page 61-7
- "Add/Edit GTP Policy Map (Details)" section on page 61-8
- "Add/Edit GTP Map" section on page 61-9

Note GTP inspection requires a special license.

GTP Inspection Overview

GPRS provides uninterrupted connectivity for mobile subscribers between GSM networks and corporate networks or the Internet. The GGSN is the interface between the GPRS wireless data network and other networks. The SGSN performs mobility, data session management, and data compression (See Figure 61-1).

Figure 61-1 GPRS Tunneling Protocol

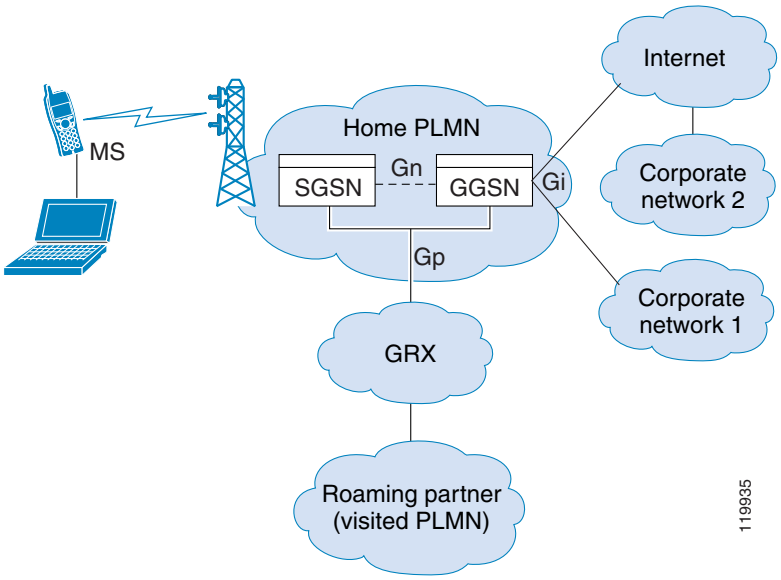

The UMTS is the commercial convergence of fixed-line telephony, mobile, Internet and computer technology. UTRAN is the networking protocol used for implementing wireless networks in this system. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN, an SGSN and the UTRAN.

GTP does not include any inherent security or encryption of user data, but using GTP with the ASA helps protect your network against these risks.

The SGSN is logically connected to a GGSN using GTP. GTP allows multiprotocol packets to be tunneled through the GPRS backbone between GSNs. GTP provides a tunnel control and management protocol that allows the SGSN to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP uses a tunneling mechanism to provide a service for carrying user data packets.

Note When using GTP with failover, if a GTP connection is established and the active unit fails before data is transmitted over the tunnel, the GTP data connection (with a "j" flag set) is not replicated to the standby unit. This occurs because the active unit does not replicate embryonic connections to the standby unit.

Select GTP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select GTP Map

The Select GTP Map dialog box lets you select or create a new GTP map. A GTP map lets you change the configuration values used for GTP application inspection. The Select GTP Map table provides a list of previously configured maps that you can select for application inspection.

Note GTP inspection requires a special license. If you try to enable GTP application inspection on a ASA without the required license, the ASA displays an error message.

Fields

- Use the default GTP inspection map—Specifies to use the default GTP map.
- Select an GTP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

GTP Inspect Map

Configuration > Global Objects > Inspect Maps > GTP

The GTP pane lets you view previously configured GTP application inspection maps. A GTP map lets you change the default configuration values used for GTP application inspection.

GTP is a relatively new protocol designed to provide security for wireless connections to TCP/IP networks, such as the Internet. You can use a GTP map to control timeout values, message sizes, tunnel counts, and GTP versions traversing the security appliance.

Note GTP inspection is not available without a special license.

Fields

- GTP Inspect Maps—Table that lists the defined GTP inspect maps.
- Add—Configures a new GTP inspect map. To edit a GTP inspect map, choose the GTP entry in the GTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the GTP Inspect Maps table.
- Security Level—Security level low only.
  - Do not Permit Errors
  - Maximum Number of Tunnels: 500
  - GSN timeout: 00:30:00
  - Pdp-Context timeout: 00:30:00
  - Request timeout: 00:01:00
  - Signaling timeout: 00:30:00.
  - Tunnel timeout: 01:00:00.
  - T3-response timeout: 00:00:20.
  - Drop and log unknown message IDs.
- IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
- Customize—Opens the Add/Edit GTP Policy Map dialog box for additional settings.

- Default Level—Sets the security level back to the default.

IMSI Prefix Filtering

Configuration > Global Objects > Inspect Maps > GTP > IMSI Prefix Filtering

The IMSI Prefix tab lets you define the IMSI prefix to allow within GTP requests.

Fields

- Mobile Country Code—Defines the non-zero, three-digit value identifying the mobile country code. One or two-digit entries will be prepended by 0 to create a three-digit value.
- Mobile Network Code—Defines the two or three-digit value identifying the network code.
- Add—Add the specified country code and network code to the IMSI Prefix table.
- Delete—Deletes the specified country code and network code from the IMSI Prefix table.

Add/Edit GTP Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Basic View

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Security level low only.

  Do not Permit Errors

  Maximum Number of Tunnels: 500

  GSN timeout: 00:30:00

  Pdp-Context timeout: 00:30:00

  Request timeout: 00:01:00

  Signaling timeout: 00:30:00.

  Tunnel timeout: 01:00:00.

  T3-response timeout: 00:00:20.

  Drop and log unknown message IDs.

  – IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
  – Default Level—Sets the security level back to the default.
- Details—Shows the Parameters, IMSI Prefix Filtering, and Inspections tabs to configure additional settings.

Add/Edit GTP Policy Map (Details)

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Advanced View

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Shows the security level and IMSI prefix filtering settings to configure.
- Permit Parameters—Tab that lets you configure the permit parameters for the GTP inspect map.
  - Object Groups to Add

    From object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.

    To object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.
  - Add—Add the specified country code and network code to the IMSI Prefix table.
  - Delete—Deletes the specified country code and network code from the IMSI Prefix table.
  - Permit Errors—Lets any packets that are invalid or that encountered an error during inspection to be sent through the ASA instead of being dropped. By default, all invalid packets or packets that failed during parsing are dropped.
- General Parameters—Tab that lets you configure the general parameters for the GTP inspect map.
  - Maximum Number of Requests—Lets you change the default for the maximum request queue size allowed. The default for the maximum request queue size is 200. Specifies the maximum number of GTP requests that will be queued waiting for a response. The permitted range is from 1 to 9999999.
  - Maximum Number of Tunnels—Lets you change the default for the maximum number of tunnels allowed. The default tunnel limit is 500. Specifies the maximum number of tunnels allowed. The permitted range is from 1 to 9999999 for the global overall tunnel limit.
  - Timeouts

    GSN timeout—Lets you change the default for the maximum period of inactivity before a GSN is removed. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

    PDP-Context timeout—Lets you change the default for the maximum period of inactivity before receiving the PDP Context for a GTP session. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

    Request Queue—Lets you change the default for the maximum period of inactivity before receiving the GTP message during a GTP session. The default is 1 minute. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

Signaling—Lets you change the default for the maximum period of inactivity before a GTP signaling is removed. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

Tunnel—Lets you change the default for the maximum period of inactivity for the GTP tunnel. The default is 1 hour. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down Request timeout—Specifies the GTP Request idle timeout.

T3-Response timeout—Specifies the maximum wait time for a response before removing the connection.

- IMSI Prefix Filtering—Tab that lets you configure the IMSI prefix filtering for the GTP inspect map.
  - Mobile Country Code—Defines the non-zero, three-digit value identifying the mobile country code. One or two-digit entries will be prepended by 0 to create a three-digit value.
  - Mobile Network Code—Defines the two or three-digit value identifying the network code.
  - Add—Add the specified country code and network code to the IMSI Prefix table.
  - Delete—Deletes the specified country code and network code from the IMSI Prefix table.
- Inspections—Tab that lets you configure the GTP inspect maps.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the GTP inspection.
  - Value—Shows the value to match in the GTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add GTP Inspect dialog box to add an GTP inspection.
  - Edit—Opens the Edit GTP Inspect dialog box to edit an GTP inspection.
  - Delete—Deletes an GTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit GTP Map

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Add/Edit GTP Map

The Add/Edit GTP Inspect dialog box lets you define the match criterion and value for the GTP inspect map.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of GTP traffic to match.
  - Access Point Name—Match on access point name.
  - Message ID—Match on the message ID.

  - Message Length—Match on the message length
  - Version—Match on the version.
- Access Point Name Criterion Values—Specifies an access point name to be matched. By default, all messages with valid APNs are inspected, and any APN is allowed.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
  - Action—Drop.
  - Log—Enable or disable.
- Message ID Criterion Values—Specifies the numeric identifier for the message that you want to match. The valid range is 1 to 255. By default, all valid message IDs are allowed.
  - Value—Specifies whether value is an exact match or a range.

    Equals—Enter a value.

    Range—Enter a range of values.
  - Action—Drop packet or limit rate (pps).
  - Log—Enable or disable.
- Message Length Criterion Values—Lets you change the default for the maximum message length for the UDP payload that is allowed.
  - Minimum value—Specifies the minimum number of bytes in the UDP payload. The range is from 1 to 65536.
  - Maximum value—Specifies the maximum number of bytes in the UDP payload. The range is from 1 to 65536.
  - Action—Drop packet.
  - Log—Enable or disable.
- Version Criterion Values—Specifies the GTP version for messages that you want to match. The valid range is 0-255. Use 0 to identify Version 0 and 1 to identify Version 1. Version 0 of GTP uses port 3386, while Version 1 uses port 2123. By default all GTP versions are allowed.
  - Value—Specifies whether value is an exact match or a range.

    Equals—Enter a value.

    Range—Enter a range of values.
  - Action—Drop packet.
  - Log—Enable or disable.

RADIUS Accounting Inspection

This section describes the IM inspection engine. This section includes the following topics:

- RADIUS Accounting Inspection Overview, page 61-11

- Select RADIUS Accounting Map, page 61-11
- Add RADIUS Accounting Policy Map, page 61-11
- RADIUS Inspect Map, page 61-12
- RADIUS Inspect Map Host, page 61-12
- RADIUS Inspect Map Other, page 61-13

RADIUS Accounting Inspection Overview

One of the well known problems is the over-billing attack in GPRS networks. The over-billing attack can cause consumers anger and frustration by being billed for services that they have not used. In this case, a malicious attacker sets up a connection to a server and obtains an IP address from the SGSN. When the attacker ends the call, the malicious server will still send packets to it, which gets dropped by the GGSN, but the connection from the server remains active. The IP address assigned to the malicious attacker gets released and reassigned to a legitimate user who will then get billed for services that the attacker will use.

RADIUS accounting inspection prevents this type of attack by ensuring the traffic seen by the GGSN is legitimate. With the RADIUS accounting feature properly configured, the security appliance tears down a connection based on matching the Framed IP attribute in the Radius Accounting Request Start message with the Radius Accounting Request Stop message. When the Stop message is seen with the matching IP address in the Framed IP attribute, the security appliance looks for all connections with the source matching the IP address.

You have the option to configure a secret pre-shared key with the RADIUS server so the security appliance can validate the message. If the shared secret is not configured, the security appliance does not need to validate the source of the message and will only check that the source IP address is one of the configured addresses allowed to send the RADIUS messages.

Note When using RADIUS accounting inspection with GPRS enabled, theASA checks for the 3GPP-Session-Stop-Indicator in the Accounting Request STOP messages to properly handle secondary PDP contexts. Specifically, the ASA requires that the Accounting Request STOP messages include the 3GPP-SGSN-Address attribute before it will temrinate the user sessions and all associated connections. Some third-party GGSNs might not send this attribute by default.

Select RADIUS Accounting Map

The Select RADIUS Accounting Map dialog box lets you select a defined RADIUS accounting map or define a new one.

Fields

- Add—Lets you add a new RADIUS accounting map.

Add RADIUS Accounting Policy Map

The Add RADIUS Accounting Policy Map dialog box lets you add the basic settings for the RADIUS accounting map.

Fields

- Name—Enter the name of the previously configured RADIUS accounting map.
- Description—Enter the description of the RADIUS accounting map, up to 100 characters in length.
- Host Parameters tab:
  - Host IP Address—Specify the IP address of the host that is sending the RADIUS messages.
  - Key: (optional)—Specify the key.
  - Add—Adds the host entry to the Host table.
  - Delete—Deletes the host entry from the Host table.
- Other Parameters tab:
  - Attribute Number—Specify the attribute number to validate when an Accounting Start is received.
  - Add—Adds the entry to the Attribute table.
  - Delete—Deletes the entry from the Attribute table.
  - Send response to the originator of the RADIUS message—Sends a message back to the host from which the RADIUS message was sent.
  - Enforce timeout—Enables the timeout for users.
  - Users Timeout—Timeout for the users in the database (hh:mm:ss).

RADIUS Inspect Map

The RADIUS pane lets you view previously configured RADIUS application inspection maps. A RADIUS map lets you change the default configuration values used for RADIUS application inspection. You can use a RADIUS map to protect against an overbilling attack.

Fields

- Name—Enter the name of the inspect map, up to 40 characters in length.
- Description—Enter the description of the inspect map, up to 200 characters in length.
- RADIUS Inspect Maps—Table that lists the defined RADIUS inspect maps. The defined inspect maps are also listed in the RADIUS area of the Inspect Maps tree.
- Add—Adds the new RADIUS inspect map to the defined list in the RADIUS Inspect Maps table and to the RADIUS area of the Inspect Maps tree. To configure the new RADIUS map, select the RADIUS entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the RADIUS Inspect Maps table and from the RADIUS area of the Inspect Maps tree.

RADIUS Inspect Map Host

The RADIUS Inspect Map Host Parameters pane lets you configure the host parameter settings for the inspect map.

Fields

- Name—Shows the name of the previously configured RADIUS accounting map.
- Description—Enter the description of the RADIUS accounting map, up to 200 characters in length.
- Host Parameters—Lets you configure host parameters.
  - Host IP Address—Specify the IP address of the host that is sending the RADIUS messages.
  - Key: (optional)—Specify the key.
- Add—Adds the host entry to the Host table.
- Delete—Deletes the host entry from the Host table.

RADIUS Inspect Map Other

The RADIUS Inspect Map Other Parameters pane lets you configure additional parameter settings for the inspect map.

Fields

- Name—Shows the name of the previously configured RADIUS accounting map.
- Description—Enter the description of the RADIUS accounting map, up to 200 characters in length.
- Other Parameters—Lets you configure additional parameters.
  - Send response to the originator of the RADIUS message—Sends a message back to the host from which the RADIUS message was sent.
  - Enforce timeout—Enables the timeout for users.

    Users Timeout—Timeout for the users in the database (hh:mm:ss).
  - Enable detection of GPRS accounting—Enables detection of GPRS accounting. This option is only available when GTP/GPRS license is enabled.
  - Validate Attribute—Attribute information.

    Attribute Number—Specify the attribute number to validate when an Accounting Start is received.

    Add—Adds the entry to the Attribute table.

    Delete—Deletes the entry from the Attribute table.

RSH Inspection

RSH inspection is enabled by default. The RSH protocol uses a TCP connection from the RSH client to the RSH server on TCP port 514. The client and server negotiate the TCP port number where the client listens for the STDERR output stream. RSH inspection supports NAT of the negotiated port number if necessary.

SNMP Inspection

This section describes the IM inspection engine. This section includes the following topics:

- SNMP Inspection Overview, page 61-14

- "Select SNMP Map" section on page 61-14
- "SNMP Inspect Map" section on page 61-14

SNMP Inspection Overview

SNMP application inspection lets you restrict SNMP traffic to a specific version of SNMP. Earlier versions of SNMP are less secure; therefore, denying certain SNMP versions may be required by your security policy. The ASA can deny SNMP versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map.

You then apply the SNMP map when you enable SNMP inspection according to the "Configuring Application Layer Protocol Inspection" section on page 57-6.

Select SNMP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select SNMP Map

The Select SNMP Map dialog box lets you select or create a new SNMP map. An SNMP map lets you change the configuration values used for SNMP application inspection. The Select SNMP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SNMP inspection map—Specifies to use the default SNMP map.
- Select an SNMP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

SNMP Inspect Map

Configuration > Global Objects > Inspect Maps > SNMP

The SNMP pane lets you view previously configured SNMP application inspection maps. An SNMP map lets you change the default configuration values used for SNMP application inspection.

Fields

- Map Name—Lists previously configured application inspection maps. Select a map and click **Edit** to view or change an existing map.
- Add—Configures a new SNMP inspect map.
- Edit—Edits the selected SNMP entry in the SNMP Inspect Maps table.
- Delete—Deletes the inspect map selected in the SNMP Inspect Maps table.

Add/Edit SNMP Map

Configuration > Global Objects > Inspect Maps > SNMP > Add/Edit SNMP Map (You can get to this dialog box through various paths.)

The Add/Edit SNMP Map dialog box lets you create a new SNMP map for controlling SNMP application inspection.

Fields

- SNMP Map Name—Defines the name of the application inspection map.
- SNMP version 1—Enables application inspection for SNMP version 1.
- SNMP version 2 (party based)—Enables application inspection for SNMP version 2.
- SNMP version 2c (community based)—Enables application inspection for SNMP version 2c.
- SNMP version 3—Enables application inspection for SNMP version 3.

XDMCP Inspection

XDMCP inspection is enabled by default; however, the XDMCP inspection engine is dependent upon proper configuration of the **established** command.

XDMCP is a protocol that uses UDP port 177 to negotiate X sessions, which use TCP when established.

For successful negotiation and start of an XWindows session, the ASA must allow the TCP back connection from the Xhosted computer. To permit the back connection, use the **established** command on the ASA. Once XDMCP negotiates the port to send the display, The **established** command is consulted to verify if this back connection should be permitted.

During the XWindows session, the manager talks to the display Xserver on the well-known port 6000 | n. Each display has a separate connection to the Xserver, as a result of the following terminal setting.

```
setenv DISPLAY Xserver:n
```

where *n* is the display number.

When XDMCP is used, the display is negotiated using IP addresses, which the ASA can NAT if needed. XDCMP inspection does not support PAT.

PART 11

Configuring Unified Communications

Information About Cisco Unified Communications Proxy Features

This chapter describes how to configure the adaptive security appliance for Cisco Unified Communications Proxy features.

This chapter includes the following sections:

- Information About the Adaptive Security Appliance in Cisco Unified Communications, page 63-1
- TLS Proxy Applications in Cisco Unified Communications, page 63-3
- Licensing for Cisco Unified Communications Proxy Features, page 63-4

Information About the Adaptive Security Appliance in Cisco Unified Communications

This section describes the Cisco UC Proxy features on the Cisco ASA 5500 series appliances. The purpose of a proxy is to terminate and reoriginate connections between a client and server. The proxy delivers a range of security functions such as traffic inspection, protocol conformance, and policy control to ensure security for the internal network. An increasingly popular function of a proxy is to terminate encrypted connections in order to apply security policies while maintaining confidentiality of connections. The Cisco ASA 5500 Series appliances are a strategic platform to provide proxy functions for unified communications deployments.

The Cisco UC Proxy includes the following solutions:

Phone Proxy: Secure remote access for Cisco encrypted endpoints, and VLAN traversal for Cisco softphones

The phone proxy feature enables termination of Cisco SRTP/TLS-encrypted endpoints for secure remote access. The phone proxy allows large scale deployments of secure phones without a large scale VPN remote access hardware deployment. End-user infrastructure is limited to just the IP endpoint, without VPN tunnels or hardware.

The Cisco adaptive security appliance phone proxy is the replacement product for the Cisco Unified Phone Proxy. Additionally, the phone proxy can be deployed for voice/data VLAN traversal for softphone applications. Cisco IP Communicator (CIPC) traffic (both media and signaling) can be proxied through the ASA, thus traversing calls securely between voice and data VLANs.

For information about the differences between the TLS proxy and phone proxy, go to the following URL for Unified Communications content, including TLS Proxy vs. Phone Proxy white paper:

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

TLS Proxy: Decryption and inspection of Cisco Unified Communications encrypted signaling

End-to-end encryption often leaves network security appliances "blind" to media and signaling traffic, which can compromise access control and threat prevention security functions. This lack of visibility can result in a lack of interoperability between the firewall functions and the encrypted voice, leaving businesses unable to satisfy both of their key security requirements.

The ASA is able to intercept and decrypt encrypted signaling from Cisco encrypted endpoints to the Cisco Unified Communications Manager (Cisco UCM), and apply the required threat protection and access control. It can also ensure confidentiality by re-encrypting the traffic onto the Cisco UCM servers.

Typically, the ASA TLS Proxy functionality is deployed in campus unified communications network. This solution is ideal for deployments that utilize end to end encryption and firewalls to protect Unified Communications Manager servers.

Mobility Proxy: Secure connectivity between Cisco Unified Mobility Advantage server and Cisco Unified Mobile Communicator clients

Cisco Unified Mobility solutions include the Cisco Unified Mobile Communicator (Cisco UMC), an easy-to-use software application for mobile handsets that extends enterprise communications applications and services to mobile phones and the Cisco Unified Mobility Advantage (Cisco UMA) server. The Cisco Unified Mobility solution streamlines the communication experience, enabling single number reach and integration of mobile endpoints into the Unified Communications infrastructure.

The security appliance acts as a proxy, terminating and reoriginating the TLS signaling between the Cisco UMC and Cisco UMA. As part of the proxy security functionality, inspection is enabled for the Cisco UMA Mobile Multiplexing Protocol (MMP), the protocol between Cisco UMC and Cisco UMA.

Presence Federation Proxy: Secure connectivity between Cisco Unified Presence servers and Cisco/Microsoft Presence servers

Cisco Unified Presence solution collects information about the availability and status of users, such as whether they are using communication devices, such as IP phones at particular times. It also collects information regarding their communications capabilities, such as whether web collaboration or video conferencing is enabled. Using user information captured by Cisco Unified Presence, applications such as Cisco Unified Personal Communicator and Cisco UCM can improve productivity by helping users connect with colleagues more efficiently through determining the most effective way for collaborative communication.

Using the ASA as a secure presence federation proxy, businesses can securely connect their Cisco Unified Presence (Cisco UP) servers to other Cisco or Microsoft Presence servers, enabling intra-enterprise communications. The security appliance terminates the TLS connectivity between the servers, and can inspect and apply policies for the SIP communications between the servers.

Cisco Intercompany Media Engine Proxy: Secure connectivity between Cisco UCM servers in different enterprises for IP Phone traffic

As more unified communications are deployed within enterprises, cases where business-to-business calls utilize unified communications on both sides with the Public Switched Network (PSTN) in the middle become increasingly common. All outside calls go over circuits to telephone providers and from there are delivered to all external destinations.

The Cisco Intercompany Media Engine gradually creates dynamic, encrypted VoIP connections between businesses, so that a collection of enterprises that work together end up looking like one giant business with secure VoIP interconnections between them.

There are three components to a Cisco Intercompany Media Engine deployment within an enterprise: a Cisco Intercompany Media Engine server, a call agent (the Cisco Unified Communications Manager) and an ASA running the Cisco Intercompany Media Engine Proxy.

The ASA provides perimeter security by encrypting signaling connections between enterprises and preventing unathorized calls. An ASA running the Cisco Intercompany Media Engine Proxy can either be deployed as an Internet firewall or be designated as a Cisco Intercompany Media Engine Proxy and placed in the DMZ, off the path of the regular Internet traffic.

TLS Proxy Applications in Cisco Unified Communications

Table 63-1 shows the Cisco Unified Communications applications that utilize the TLS proxy on the ASA.

Table 63-1 TLS Proxy Applications and the Security Appliance

Application	TLS Client	TLS Server	Client Authentication	Security Appliance Server Role	Security Appliance Client Role
Phone Proxy and TLS Proxy	IP phone	Cisco UCM	Yes	Proxy certificate, self-signed or by internal CA	Local dynamic certificate signed by the ASA CA (might not need certificate for phone proxy application)
Mobility Proxy	Cisco UMC	Cisco UMA	No	Using the Cisco UMA private key or certificate impersonation	Any static configured certificate
Presence Federation Proxy	Cisco UP or MS LCS/OCS	Cisco UP or MS LCS/OCS	Yes	Proxy certificate, self-signed or by internal CA	Using the Cisco UP private key or certificate impersonation

The ASA supports TLS proxy for various voice applications. For the phone proxy, the TLS proxy running on the ASA has the following key features:

- The ASA forces remote IP phones connecting to the phone proxy through the Internet to be in secured mode even when the Cisco UCM cluster is in non-secure mode.
- The TLS proxy is implemented on the ASA to intercept the TLS signaling from IP phones.
- The TLS proxy decrypts the packets, sends packets to the inspection engine for NAT rewrite and protocol conformance, optionally encrypts packets, and sends them to Cisco UCM or sends them in clear text if the IP phone is configured to be in nonsecure mode on the Cisco UCM.
- The ASA acts as a media terminator as needed and translates between SRTP and RTP media streams.
- The TLS proxy is a transparent proxy that works based on establishing trusted relationship between the TLS client, the proxy (the ASA), and the TLS server.

For the Cisco Unified Mobility solution, the TLS client is a Cisco UMA client and the TLS server is a Cisco UMA server. The ASA is between a Cisco UMA client and a Cisco UMA server. The mobility proxy (implemented as a TLS proxy) for Cisco Unified Mobility allows the use of an imported PKCS-12 certificate for server proxy during the handshake with the client. Cisco UMA clients are not required to present a certificate (no client authentication) during the handshake.

For the Cisco Unified Presence solution, the ASA acts as a TLS proxy between the Cisco UP server and the foreign server. This allows the ASA to proxy TLS messages on behalf of the server that initiates the TLS connection, and route the proxied TLS messages to the client. The ASA stores certificate trustpoints for the server and the client, and presents these certificates on establishment of the TLS session.

Licensing for Cisco Unified Communications Proxy Features

The Cisco Unified Communications proxy features supported by the ASA require a Unified Communications Proxy license:

- Phone proxy
- TLS proxy for encrypted voice inspection
- Presence federation proxy
- Intercompany media engine proxy

Note In Version 8.2(2) and later, the Mobility Advantage proxy no longer requires a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform for the phone proxy, TLS proxy for encrypted voice inspection, and presence federation proxy:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5510	Base License and Security Plus License: 2 sessions. *Optional licenses: 24, 50, or 100 sessions.*
ASA 5520	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5540	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5550	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5580	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]

Model	License Requirement[1]
ASA 5512-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
   - Phone Proxy
   - Presence Federation Proxy
   - Encrypted Voice Inspection

   Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

   Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

   You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

   Note: For license part numbers ending in “K8” (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in “K9” (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

   Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

   You might also use SRTP encryption sessions for your connections:
   - For K8 licenses, SRTP sessions are limited to 250.
   - For K9 licenses, there is not limit.

   Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.
2. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

Table 63-2 shows the default and maximum TLS session details by platform.

Table 63-2 Default and Maximum TLS Sessions on the Security Appliance

Security Appliance Platform	Default TLS Sessions	Maximum TLS Sessions
ASA 5505	10	80
ASA 5510	100	200
ASA 5520	300	1200
ASA 5540	1000	4500
ASA 5550	2000	4500
ASA 5580	4000	13,000

The following table shows the Unified Communications Proxy license details by platform for intercompany media engine proxy:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	Intercompany Media Engine license. When you enable the Intercompany Media Engine (IME) license, you can use TLS proxy sessions up to the configured TLS proxy limit. If you also have a Unified Communications (UC) license installed that is higher than the default TLS proxy limit, then the ASA sets the limit to be the UC license limit plus an additional number of sessions depending on your model. You can manually configure the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. If you also install the UC license, then the TLS proxy sessions available for UC are also available for IME sessions. For example, if the configured limit is 1000 TLS proxy sessions, and you purchase a 750-session UC license, then the first 250 IME sessions do not affect the sessions available for UC. If you need more than 250 sessions for IME, then the remaining 750 sessions of the platform limit are used on a first-come, first-served basis by UC and IME. • For a license part number ending in "K8", TLS proxy sessions are limited to 1000. • For a license part number ending in "K9", the TLS proxy limit depends on your configuration and the platform model. **Note** K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted. You might also use SRTP encryption sessions for your connections: • For a K8 license, SRTP sessions are limited to 250. • For a K9 license, there is no limit. **Note** Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

For more information about licensing, see Chapter 5, "Managing Feature Licenses."

Using the Cisco Unified Communication Wizard

This chapter describes how to configure the adaptive security appliance for Cisco Unified Communications Proxy features.

This chapter includes the following sections:

- Information about the Cisco Unified Communication Wizard, page 64-1
- Licensing Requirements for the Unified Communication Wizard, page 64-3
- Guidelines and Limitations, page 64-4
- Configuring the Phone Proxy by using the Unified Communication Wizard, page 64-4
- Configuring the Mobility Advantage by using the Unified Communication Wizard, page 64-11
- Configuring the Presence Federation Proxy by using the Unified Communication Wizard, page 64-14
- Configuring the UC-IME by using the Unified Communication Wizard, page 64-16
- Working with Certificates in the Unified Communication Wizard, page 64-23

Information about the Cisco Unified Communication Wizard

Note The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

The Unified Communication Wizard assists you in configuring the following Unified Communications proxies on the ASA:

- Cisco Phone Proxy

  See Configuring the Phone Proxy by using the Unified Communication Wizard, page 64-4.
- Cisco Mobility Advantage Proxy

  See Configuring the Mobility Advantage by using the Unified Communication Wizard, page 64-11.
- Cisco Presence Federation Proxy

  See Configuring the Presence Federation Proxy by using the Unified Communication Wizard, page 64-14.
- Cisco Intercompany Media Engine Proxy

  See Configuring the UC-IME by using the Unified Communication Wizard, page 64-16.

The wizard simplifies the configuration of the Unified Communications proxies in the following ways:

- You enter all required data in the wizard steps. You are not required to navigate various ASDM screens to configure the Unified Communications proxies.
- The wizard generates configuration settings for the Unified Communications proxies where possible, automatically, without requiring you to enter data. For example, the wizard configures the required access lists, IP address translation (NAT and PAT) statements, self-signed certificates, TLS proxies, and application inspection.
- The wizard displays network diagrams to illustrate data collection.

To access the Unified Communication Wizard, choose one of the following paths in the main ASDM application window:

- **Wizards > Unified Communication Wizard**.
- **Configuration > Firewall > Unified Communications,** and then click **Unified Communication Wizard**.

Phone Proxy: Secure remote access for Cisco encrypted endpoints, and VLAN traversal for Cisco softphones

The phone proxy feature enables termination of Cisco SRTP/TLS-encrypted endpoints for secure remote access. The phone proxy allows large scale deployments of secure phones without a large scale VPN remote access hardware deployment. End-user infrastructure is limited to just the IP endpoint, without VPN tunnels or hardware.

The Cisco adaptive security appliance phone proxy is the replacement product for the Cisco Unified Phone Proxy. Additionally, the phone proxy can be deployed for voice/data VLAN traversal for softphone applications. Cisco IP Communicator (CIPC) traffic (both media and signaling) can be proxied through the ASA, thus traversing calls securely between voice and data VLANs.

For information about the differences between the TLS proxy and phone proxy, go to the following URL for Unified Communications content, including TLS Proxy vs. Phone Proxy white paper:

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

Mobility Advantage Proxy: Secure connectivity between Cisco Mobility Advantage server and Cisco Unified Mobile Communicator clients

Cisco Mobility Advantage solutions include the Cisco Unified Mobile Communicator (Cisco UMC), an easy-to-use software application for mobile handsets that extends enterprise communications applications and services to mobile phones and the Cisco Unified Mobility Advantage (Cisco UMA) server. The Cisco Mobility Advantage solution streamlines the communication experience, enabling single number reach and integration of mobile endpoints into the Unified Communications infrastructure.

The security appliance acts as a proxy, terminating and reoriginating the TLS signaling between the Cisco UMC and Cisco UMA. As part of the proxy security functionality, inspection is enabled for the Cisco UMA Mobile Multiplexing Protocol (MMP), the protocol between Cisco UMC and Cisco UMA.

Presence Federation Proxy: Secure connectivity between Cisco Unified Presence servers and Cisco/Microsoft Presence servers

Cisco Unified Presence solution collects information about the availability and status of users, such as whether they are using communication devices, such as IP phones at particular times. It also collects information regarding their communications capabilities, such as whether web collaboration or video conferencing is enabled. Using user information captured by Cisco Unified Presence, applications such as Cisco Unified Personal Communicator and Cisco UCM can improve productivity by helping users connect with colleagues more efficiently through determining the most effective way for collaborative communication.

Using the ASA as a secure presence federation proxy, businesses can securely connect their Cisco Unified Presence (Cisco UP) servers to other Cisco or Microsoft Presence servers, enabling intra-enterprise communications. The security appliance terminates the TLS connectivity between the servers, and can inspect and apply policies for the SIP communications between the servers.

Cisco Intercompany Media Engine Proxy: Secure connectivity between Cisco UCM servers in different enterprises for IP Phone traffic

As more unified communications are deployed within enterprises, cases where business-to-business calls utilize unified communications on both sides with the Public Switched Network (PSTN) in the middle become increasingly common. All outside calls go over circuits to telephone providers and from there are delivered to all external destinations.

The Cisco Intercompany Media Engine (UC-IME) gradually creates dynamic, encrypted VoIP connections between businesses, so that a collection of enterprises that work together end up looking like one giant business with secure VoIP interconnections between them.

There are three components to a Cisco Intercompany Media Engine deployment within an enterprise: a Cisco Intercompany Media Engine server, a call agent (the Cisco Unified Communications Manager) and an ASA running the Cisco Intercompany Media Engine Proxy.

The ASA provides perimeter security by encrypting signaling connections between enterprises and preventing unauthorized calls. An ASA running the Cisco Intercompany Media Engine Proxy can either be deployed as an Internet firewall or be designated as a Cisco Intercompany Media Engine Proxy and placed in the DMZ, off the path of the regular Internet traffic.

Licensing Requirements for the Unified Communication Wizard

To run the Unified Communication Wizard in ASDM, you require the following license:

Model	License Requirement
All models	Base License

However, to run each of the Unified Communications proxy features created by the wizard, you must have the appropriate Unified Communications Proxy licenses.

The Cisco Unified Communications proxy features supported by the ASA require a Unified Communications Proxy license:

- Cisco Phone Proxy
- TLS proxy for encrypted voice inspection
- Presence Federation Proxy
- Cisco Intercompany Media Engine Proxy

See Licensing for Cisco Unified Communications Proxy Features, page 63-4 for more information.

Note The Cisco Intercompany Media Engine Proxy does not appear as an option in the Unified Communication Wizard unless the license required for this proxy is installed on the ASA.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6 addresses.

Additional Guidelines and Limitations

Using the Unified Communication Wizard to create the Unified Communications proxies has the following limitations and requirements:

- You must configure at least two interfaces on the ASA to use the UC Wizard to configure a Unified Communications proxy.
- For all Unified Communications proxies to function correctly, you must synchronize the clock on the ASA and all servers associated with each proxy, such as the Cisco Unified Communication Manager server, the Cisco Mobility Advantage server, the Cisco Unified Presence server, and the Cisco Intercompany Media Engine server.
- When you configure the Cisco Intercompany Media Engine Proxy for an off-path deployment, you must ensure that the public IP addresses and ports of the Cisco Unified Communications Manager servers and the public IP address for the media termination address are accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of the requirements.
- If the ASA on which you configure the Cisco Mobility Advantage Proxy and the Cisco Presence Federation Proxy is located behind another firewall, you must ensure that the public IP addresses for the Cisco Mobility Advantage server and the Cisco Unified Presence server are accessible from the Internet.
- If you use the Unified Communication Wizard to create to the Presence Federation Proxy and the Cisco Intercompany Media Engine Proxy, you might be required to adjust the configuration of the access lists created automatically by the wizard for each proxy. See Chapter 68, “Configuring Cisco Unified Presence” and Chapter 69, “Configuring Cisco Intercompany Media Engine Proxy”, respectively, for information about the access list requirements required by each proxy.

Configuring the Phone Proxy by using the Unified Communication Wizard

To configure the Cisco Unified Presence proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Phone Proxy option under the Remote Access section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Phone Proxy instance, importing and installing the required certificates, and finally enables the SIP and SCCP inspection for the Phone Proxy traffic automatically.

Note Any configuration created by the wizard should be maintained through the wizard to ensure proper synchronization. For example, if you create a phone proxy configuration through the UC wizard and then modify the configuration outside of the wizard, the rest of the wizard configuration is not updated, and the wizard configuration is not synchronized.

Therefore, if you choose to change some part of the phone proxy configuration outside of the wizard, it is your responsibility to keep the rest of the configuration in synchronization.

The wizard guides you through four steps to configure the Phone Proxy:

Step 1 Select the Phone Proxy option.

Step 2 Specify settings to define the Cisco Unified Communications Manager (UCM) servers and TFTP servers, such the IP address and the address translation settings of each server, and the Cisco UCM cluster security mode. See Configuring the Private Network for the Phone Proxy, page 64-5 and Configuring Servers for the Phone Proxy, page 64-6.

Step 3 If required, enable Certificate Authority Proxy Function (CAPF). See Enabling Certificate Authority Proxy Function (CAPF) for IP Phones, page 64-8.

Step 4 Configure the public IP phone network, such as address translation settings for remote IP phones, whether to enable service setting for IP phones, and the HTTP proxy used by the IP phones. Configuring the Public IP Phone Network, page 64-9

Step 5 Specify the media termination address settings of the Cisco UCM. Configuring the Media Termination Address for Unified Communication Proxies, page 64-10.

The wizard completes by displaying a summary of the configuration created for Phone Proxy.

Configuring the Private Network for the Phone Proxy

The values that you specify in this page configure the connection from the ASA to the Cisco UCMs and TFTP servers by creating the necessary address translation settings and access control list entries.

Additionally, you specify the security mode for the Cisco UCM cluster. In a nonsecure cluster mode or a mixed mode where the phones are configured as nonsecure, the phone proxy behaves in the following ways:

- The TLS connections from the phones are terminated on the ASA and a TCP connection is initiated to the Cisco UCM.
- SRTP sent from external IP phones to the internal network IP phone via the ASA is converted to RTP.

In a mixed mode cluster where the internal IP phones are configured as authenticated, the TLS connection is not converted to TCP to the Cisco UCM but the SRTP is converted to RTP.

In a mixed mode cluster where the internal IP phone is configured as encrypted, the TLS connection remains a TLS connection to the Cisco UCM and the SRTP from the remote phone remains SRTP to the internal IP phone.

Step 1 From the Interface drop-down list, choose the interface on which the ASA listens for the Cisco UCM servers and TFTP servers. The Cisco UCM servers and TFTP servers must reside on the same interface.

Step 2 Specify each entity in the network (all Cisco UCM and TFTP servers) that the IP phones must trust. Click **Add** to add the servers. See Configuring Servers for the Phone Proxy, page 64-6.

To modify the configuration of a server already added to the configuration, select the server in the table and click **Edit**. The Edit Server dialog appears. See Configuring Servers for the Phone Proxy, page 64-6. At least one Cisco UCM and at least one TFTP server must be configured for the phone proxy.

Step 3 Specify the security mode of the Cisco UCM cluster by clicking one of the following options in the Unified CM Cluster Mode field:

- Non-secure—Specifies the cluster to be in nonsecure mode when configuring the Phone Proxy feature.
- Mixed—Specifies the cluster to be in mixed mode when configuring the Phone Proxy feature.

  If you selected the Mixed security mode, the Generate and Export LDC Certificate button becomes available.

Step 4 For a Mixed security mode only, configure local dynamic certificates (LDC) for the IP phones by performing the following steps:

a. Click the **Generate and Export LDC Certificate** button.

A dialog box appears stating “Enrollment succeeded,” which indicates that the LDC was generated.

b. Click **OK** to close the Enrollment Status dialog box. The Export certificate dialog box appears.

c. In the Export to File field, enter the file name and path for the LDC or click browse to locate and select an existing file.

d. Click the **Export Certificate** button. A dialog box appears indicating that the file was exported successfully.

e. Click **OK** to close the dialog box. A dialog box appears reminding you to install the LDC on the Cisco UCMs.

f. Click **OK** to close the dialog box.

Once configured, the ASA presents this unique, dynamically-created certificate to the Cisco UCM on behalf of the IP phones.

Step 5 Click **Next**.

Configuring Servers for the Phone Proxy

The values that you specify in this page generate address translation settings, access list entries, trustpoints, and the corresponding CTL file entries for each server.

You must add a server for each entity in the network that the IP phones must trust. These servers include all Cisco UCM servers in the cluster and all the TFTP servers.

You must add at least one TFTP server and at least one Cisco UCM server for the phone proxy. You can configure up to five TFTP servers for the phone proxy. The TFTP server is assumed to be behind the firewall on the trusted network; therefore, the phone proxy intercepts the requests between the IP phones and TFTP server.

The servers that the IP phones must trust can be deployed on the network in one of the following ways:

- All the services required by the Cisco UCM server, namely the Cisco UCM, TFTP, and CAPF services, are running on one server. In this deployment, only one instance of each service exists. For this deployment, you can select Unified CM+ TFTP as the server type. You can either use Address only or Address and ports for address translation. Cisco recommends that you specify Address and ports for increased security.
- Deployments for larger enterprises might have redundant Cisco UCMs and dedicated servers for TFTP and CAPF services. In that type of deployment, use Address only for voice address translation and Address only or Address and ports for TFTP.

Table 64-1 lists the ports that are configured for Address and port translation by default:

Table 64-1 Port Configuration

Address	Default Port	Description
TFTP Server	69	Allows incoming TFTP
Cisco UCM	2000	Allows incoming non-secure SCCP
Cisco UCM	2443	Allows incoming secure SCCP
Cisco UCM	5061	Allows incoming secure SIP

Step 1 In the Server Type field, select the server from the drop-down list: Unified CM, TFTP, or Unified CM + TFTP. Select Unified CM + TFTP when the Cisco UCM and TFTP server reside on the same device.

Note Depending on which type of server you select (Unified CM or TFTP), only the necessary fields in this dialog box become available. Specifically, if the server type is Unified CM, the TFTP section in the dialog is unavailable. If the server type is TFTP, the Voice section is unavailable.

Step 2 In the Private Address field, specify the actual internal IP address of the server.

Step 3 In the FQDN field, enter the fully-qualified domain name of the server, which includes the hostname and domain name; for example, `ucm.cisco.com` (where `ucm` is the hostname and `cisco.com` is the domain name).

If you are configuring a Unified CM server, enter the fully-qualified domain name configured on the Cisco UCM.

If you are configuring a TFTP server, only specify the TFTP server fully-qualified domain name when that server is configured with FQDN. If the TFTP server is not configured with FQDN, you can leave the field blank.

Note Entering the fully-qualified domain name allows the ASA to perform hostname resolution when DNS lookup is not configured on the ASA or the configured DNS servers are unavailable.See the command reference for information about the **dns domain-lookup** command.

Step 4 In the Address Translation section, select whether to use the interface IP address or to enter a different IP address.

Selecting the Use interface IP radio button configures the server to use the IP address of the public interface. You select the public interface in step 4 of the wizard when you configure the public network for the phone proxy.

If the Use interface IP radio button is selected, you must specify port translation settings in the Voice and TFTP sections. Address-only translation is available only when you specify an IP address other than the IP address of the public interface.

When you select the Address only radio button, the ASA performs address translation on all traffic between the server and the IP phones. Selecting the Address and ports radio button limits address translation to the specified ports.

Step 5 (Unified CM or Unified CM + TFTP servers only) In the Voice section, configure inspection of SIP or SCCP protocol traffic, or both SIP and SCCP protocol traffic by completing the following fields:

a. In the Translation Type field, specify whether to use the Address only or the Address and ports.

When the deployment has redundant Cisco UCM servers and dedicated servers for TFTP and CAPF services, select Address only for voice address translation.

Select the Address and ports option when you want to limit address translation to the specified ports.

b. In the Voice Protocols field, select the inspection protocols supported by the IP phones deployed in the enterprise. Depending on which inspection protocols you select—SCCP, SIP, or SCCP and SIP—only the ports fields for the selected voice protocols are available.

c. In the Port Translation section, enter the private and public ports for the voice protocols.

The default values for the voice ports appear in the text fields. If necessary, change the private ports to match the settings on the Cisco UCM. The values you set for the public ports are used by the IP phones to traverse the ASA and communicate with the Cisco UCM.

The secure SCCP private port and public port are automatically configured. These port numbers are automatically set to the value of the non-secure port number plus 443.

Step 6 (TFTP or Unified CM + TFTP servers only) In the TFTP section, you can select either Address only or Address and port for address translation. Cisco recommends that you specify Address and port for increased security. Specifying Address and port configures the TFTP server to listen on port 69 for TFTP requests.

When the server type is Unified CM + TFTP, the wizard configures the same type of address translation for Voice and TFTP; for example, when the server type is Unified CM + TFTP and the Address only option is selected, the wizard creates a global address translation rule for all traffic to and from the server. In this case, configuring port translation for the TFTP server would be redundant.

Step 7 Click **OK** to add the server to the phone proxy configuration and return to step 2 of the wizard.

Enabling Certificate Authority Proxy Function (CAPF) for IP Phones

As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via locally significant certificate (LSC) provisioning. With LSC provisioning, you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

If your network includes Cisco IP Communicators (CIPC) or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. The certificate will be used to generate the LSC on the IP phones.

If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA. However, the wizard supports configuring only one CAPF certificate, which is the default. To import more than one CAPF certificate, go to Configuration > Device Management > Certificate Management > Identity Certificates.

You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

Step 1 Check the **Enable Certificate Authority Proxy Function** check box. The remaining fields in the page become available.

Step 2 Enter the private IP address of the LSC provider.

Step 3 In the Public Address field, specify whether to use the IP address of the ASA public interface or enter an IP address.

Specifying the private and public IP addresses for the LSC provider, creates an access list entry that allows the IP phones to contact the Cisco UCM by opening the CAPF port for LSC provisioning.

Step 4 In the Translation Type field, select the Address only or Address and ports radio button.

The IP phones must contact the CAPF service on the Cisco UCM. The address translation type (Address only versus Address and ports) you select for CAPF must match the address translation type of the Cisco UCM on which the CAPF service is running. You set the address translation type for that Cisco UCM server in the previous step of this wizard (see Configuring Servers for the Phone Proxy, page 64-6),

By default, the CAPF Service uses port 3804. Modify this default value only when it is modified on the Cisco UCM.

Step 5 If you selected the Address and ports radio button, enter the private and public ports for the CAPF service.

Step 6 Click the **Install CAPF Certificate** button. The Install Certificate dialog box appears. See Installing a Certificate, page 64-23.

Step 7 Click **Next**.

Configuring the Public IP Phone Network

The values that you specify in this page generate the address translation rules used for the IP phones and configure how the ASA handles IP phone settings.

Step 1 From the Interface drop-down list, choose the interface on which the ASA listens for connections from IP phones.

Step 2 To preserve Call Manager configuration on the IP phones, check the Preserve the Unified CM's configuration on the phone's service check box. When this check box is uncheck, the following service settings are disabled on the IP phones:

- Web Access
- PC Port
- Voice VLAN access

- Gratuitous ARP
- Span to PC Port

Step 3 To configure address translation for IP phones, check the Enable address translation for IP phones check box. Select whether to use the IP address of the ASA private interface (which you selected in step 2 of the wizard) or enter an IP address.

Configuring address translation for IP phone configures the address used by the IP phones. All traffic from the outside network converges into one source IP address so that, if there is another corporate firewall in the network, a pinhole needs to be opened only for that IP address rather than for all traffic.

Step 4 To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the <proxyServerURL> tag, do the following:

a. Check the Configure an HTTP proxy to redirect phone URLs... check box.

b. In the IP Address field, type the IP address of the HTTP proxy

c. In the Port field, enter the listening port of the HTTP proxy.

The IP address you enter should be the global IP address based on where the IP phone and HTTP proxy server is located. You can enter a hostname in the IP Address field when that hostname can be resolved to an IP address by the adaptive security appliance (for example, DNS lookup is configured) because the adaptive security appliance will resolve the hostname to an IP address. If a port is not specified, the default will be 8080.

d. In the Interface field, select the interface on which the HTTP proxy resides on the adaptive security appliance.

Setting the proxy server configuration option for the Phone Proxy allows for an HTTP proxy on the DMZ or external network in which all the IP phone URLs are directed to the proxy server for services on the phones. This setting accommodates nonsecure HTTP traffic, which is not allowed back into the corporate network.

Step 5 Click **Next**.

Configuring the Media Termination Address for Unified Communication Proxies

The data from this step generates the MTA instance to be added to the Phone Proxy and the UC-IME proxy.

The phone proxy and the UC-IME proxy use the media termination address for Secure RTP (SRTP) and RTP traffic. SRTP traffic sent from external IP phones to the internal network IP phone via the ASA is converted to RTP traffic. The traffic is terminated on the adaptive security appliance. SRTP provides message authentication and replay protection to Internet media traffic such as audio and video. RTP defines a standardized packet format for delivering audio and video over the Internet.

For the UC-IME proxy and the Phone Proxy to be fully functional, you must ensure that the public IP address for the media termination address (MTA) is accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of this requirement.

The MTA IP addresses that you specify must meet specific requirements. See Media Termination Instance Prerequisites, page 65-6 for information.

Step 1 In the field for the private IP address, enter the IP address on which private media traffic terminates. The IP address must be within the same subnet as the private interface IP address. The correct subnet range is provided to the right of the field for the private IP address.

Step 2 In the field for the public IP address, enter the IP address on which public media traffic terminates. The IP address must be within the same subnet as the public interface IP address. The correct subnet range is provided to the right of the field for the public IP address.

Step 3 Specify the minimum and maximum values for the RTP port range for the media termination instance.

Port values must be within the range of 1024 to 65535.

Step 4 Click **Next**.

The wizard completes by displaying a summary of the configuration created for proxy.

Configuring the Mobility Advantage by using the Unified Communication Wizard

Note The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

The Unified Communication wizard guides you through the steps to configure the Mobility Advantage proxy. Choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. Click the Cisco Mobility Advantage Proxy radio button under the Remote Access section.

When using the wizard to create the Mobility Advantage proxy, ASDM automatically creates the necessary TLS proxies, enables MMP inspection for the Mobility Advantage traffic, generates address translation (NAT) statements, and creates the access rules that are necessary to allow traffic between the Cisco Mobility Advantage server and the mobility clients.

The following steps provide the high-level overview for configuring the Mobility Advantage proxy:

Step 1 Specify settings to define the private and public network topology, such the public and private network interfaces, and the IP addresses of the Cisco Mobility Advantage server. See Configuring the Topology for the Cisco Mobility Advantage Proxy, page 64-12.

Step 2 Configure the certificates that are exchanged between the Cisco Mobility Advantage server and the ASA. See Configuring the Server-Side Certificates for the Cisco Mobility Advantage Proxy, page 64-12.

Step 3 Configure the client-side certificate management, namely the certificates that are exchanged between the Unified Mobile Communicator clients and the ASA. See Configuring the Client-Side Certificates for the Cisco Mobility Advantage Proxy, page 64-13.

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy.

Configuring the Topology for the Cisco Mobility Advantage Proxy

When configuring the Mobility Advantage Proxy, you specify settings to define the private and public network topology, such the private and public network interfaces, and the private and public IP addresses of the Cisco Mobility Advantage server.

The values that you specify in this page generate the following configuration settings for the Mobility Advantage Proxy:

- Static PAT for the Cisco Mobility Advantage server
- Static NAT for Cisco Unified Mobile Communicator clients if the Enable address translation for Mobility clients check box is checked.
- Access lists to allow Cisco Unified Mobile Communicator clients to access the Cisco Mobility Advantage server

Step 1 In the Private Network area, choose the interface from the drop-down list.

Step 2 In the Unified MA Server area, enter the private and public IP address for the Cisco Mobility Advantage server. Entering ports for these IP addresses is optional. By default port number 5443 is entered, which is the default TCP port for MMP inspection.

Step 3 In the FQDN field, enter the domain name for the Cisco Mobility Advantage server. This domain name is included in the certificate signing request that you generate later in this wizard.

Step 4 In the Public Network area, choose an interface from the drop-down list.

The proxy uses this interface for configuring static PAT for the Cisco Mobility Advantage server and the access lists to allow Cisco Unified Mobile Communicator clients to access the Cisco Mobility Advantage server.

Step 5 **To configure whether address translation (NAT) is used by** Cisco Unified Mobile Communicator clients, check the **Enable address translation for Mobility clients** check box and choose whether to use the IP address of the public interface or whether to enter an IP address.

Step 6 Click **Next**.

Configuring the Server-Side Certificates for the Cisco Mobility Advantage Proxy

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The ASA's identity certificate is exported, and then uploaded on the Cisco UMA server truststore. The Cisco UMA server certificate is downloaded, and then uploaded on the ASA truststore.

The supports using self-signed certificates only at this step.

Step 1 In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog boxes appear indicating that the enrollment seceded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

Note

- If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.

- When using the wizard to configure the Cisco Mobility Advantage proxy, the wizard only supports installing self-signed certificates.

Step 2 Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 64-23.

Step 3 In the Unified MA Server's Certificate area, click **Install Unified MA Server's Certificate**. The Install Certificate dialog appears.

Step 4 Locate the file containing the Cisco Mobility Advantage server certificate or paste the certificate details in the dialog box. See Installing a Certificate, page 64-23.

Step 5 Click **Next**.

Note See the Cisco Mobility Advantage server documentation for information on how to export the certificate for this server.

Configuring the Client-Side Certificates for the Cisco Mobility Advantage Proxy

To establish a trust relationship between the Cisco Unified Mobile Communicator (UMC) clients and the ASA, the ASA uses a CA-signed certificate that is configured with the Cisco Mobility Advantage server's FQDN (also referred to as certificate impersonation).

In the Client-Side Certificate Management page, you enter both the intermediate CA certificate (if applicable, as in the cases of Verisign) and the signed ASA identity certificate.

Note If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 2.

Step 1 In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 64-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 64-25.

Step 2 Click **Install ASA's Identity Certificate**. Install the certificate. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 64-26.

Step 3 Click **Install Root CA's Certificate**. The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 64-23.

Step 4 Click **Next**.

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy.

Configuring the Presence Federation Proxy by using the Unified Communication Wizard

Note The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

To configure the Cisco Unified Presence proxy by using ASDM, choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. From the first page, select the Cisco Unified Presence Proxy option under the Business-to-Business section.

When using the wizard to create the Cisco Presence Federation proxy, ASDM automatically creates the necessary TLS proxies, enables SIP inspection for the Presence Federation traffic, generates address translation (static PAT) statements for the local Cisco Unified Presence server, and creates access lists to allow traffic between the local Cisco Unified Presence server and remote servers.

The following steps provide the high-level overview for configuring the Presence Federation Proxy:

Step 1 Specify settings to define the private and public network topology, such the private and public IP address of the Presence Federation server. See Configuring the Topology for the Cisco Presence Federation Proxy, page 64-14.

Step 2 Configure the local-side certificate management, namely the certificates that are exchanged between the local Unified Presence Federation server and the ASA. See Configuring the Local-Side Certificates for the Cisco Presence Federation Proxy, page 64-15.

Step 3 Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. See Configuring the Remote-Side Certificates for the Cisco Presence Federation Proxy, page 64-15.

The wizard completes by displaying a summary of the configuration created for the Presence Federation proxy.

Configuring the Topology for the Cisco Presence Federation Proxy

When configuring the Presence Federation Proxy, you specify settings to define the private and public network topology, such the private and public network interfaces, and the private and public IP addresses of the Cisco Unified Presence server.

The values that you specify in this page generate the following configuration settings for the Presence Federation Proxy:

- Static PAT for the local Cisco Unified Presence server
- Access lists for traffic between the local Cisco Unified Presence server and remote servers

Step 1 In the Private Network area, choose the interface from the drop-down list.

Step 2 In the Unified Presence Server area, enter the private and public IP address for the Unified Presence server. Entering ports for these IP addresses is optional. By default port number 5061 is entered, which is the default TCP port for SIP inspection.

Step 3 In the FQDN field, enter the domain name for the Unified Presence server. This domain name is included in the certificate signing request that you generate later in this wizard.

Step 4 In the Public Network area, choose the interface of the public network from the drop-down list. The proxy uses this interface for configuring static PAT for the local Cisco Unified Presence server and for configuring access lists to allow remote servers to access the Cisco Unified Presence server.

Step 5 Click **Next**.

Configuring the Local-Side Certificates for the Cisco Presence Federation Proxy

Within an enterprise, setting up a trust relationship is achievable by using self-signed certificates. The supports using self-signed certificates only at this step.

Step 1 In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog box appears indicating that enrollment succeeded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

Note

- If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.
- When using the wizard to configure the Cisco Presence Federation proxy, the wizard only supports installing self-signed certificates.

Step 2 Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 64-23.

Step 3 Local Unified Presence Server's Certificate area, click **Install Server's Certificate**. The Install Certificate dialog appears.

Step 4 Locate the file containing the Cisco Unified Presence server certificate or paste the certificate details in the dialog box. See Installing a Certificate, page 64-23.

Step 5 Click **Next**.

Note See the Cisco Unified Presence server documentation for information on how to export the certificate for this server.

Configuring the Remote-Side Certificates for the Cisco Presence Federation Proxy

Establishing a trust relationship across enterprises or across administrative domains is key for federation. Across enterprises you must use a trusted third-party CA (such as, VeriSign). The security appliance obtains a certificate with the FQDN of the Cisco Unified Presence server (certificate impersonation).

For the TLS handshake, the two entities, namely the local entity and a remote entity, could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. The local entity and the remote entity enroll with the CAs. The ASA as the TLS proxy must be trusted by both the local and remote entities. The security appliance is always associated with one of the enterprises. Within that enterprise, the entity and the security appliance authenticate each other by using a self-signed certificate.

To establish a trusted relationship between the security appliance and the remote entity, the security appliance can enroll with the CA on behalf of the Cisco Unified Presence server for the local entity. In the enrollment request, the local entity identity (domain name) is used.

To establish the trust relationship, the security appliance enrolls with the third party CA by using the Cisco Unified Presence server FQDN as if the security appliance is the Cisco Unified Presence server.

Note If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 2.

Step 1 In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 64-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 64-25.

Step 2 Click **Install ASA's Identity Certificate**. See Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 64-26.

Step 3 Click **Remote Server's CA's Certificate**. The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 64-23.

Note You must install a root CA certificate for each remote entity that communicates with the ASA because different organizations might be using different CAs.

Step 4 Click **Next**.

The wizard completes by displaying a summary of the configuration created for the Presence Federation proxy.

Configuring the UC-IME by using the Unified Communication Wizard

Note The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

To configure the Cisco Intercompany Media Engine Proxy by using ASDM, choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. From the first page, select the Cisco Intercompany Media Engine Proxy option under the Business-to-Business section and click **Next**.

Note The Cisco Intercompany Media Engine Proxy does not appear as an option in the Unified Communication Wizard unless the license required for this proxy is installed on the ASA.

When using the wizard to create the Cisco Intercompany Media Engine Proxy, ASDM automatically creates the necessary TLS proxies, enables SIP inspection for Cisco Intercompany Media Engine traffic, generates address translation (static PAT) statements for local Cisco Unified Communications Manager servers, and creates access lists to allow traffic between the local Cisco Unified Communications Manager servers and the remote servers.

The following steps provide the high-level overview for configuring the Cisco Intercompany Media Engine Proxy:

Step 1 Select the topology of the Cisco Intercompany Media Engine Proxy, namely whether the security appliance is an edge firewall with all Internet traffic flowing through it or whether the security appliance is off the path of the main Internet traffic (referred to as an off-path deployment). See Configuring the Topology for the Cisco Intercompany Media Engine Proxy, page 64-17.

Step 2 Specify private network settings such as the Cisco UCM IP addresses and the ticket settings. See Configuring the Private Network Settings for the Cisco Intercompany Media Engine Proxy, page 64-18.

Step 3 Specify the public network settings. See Configuring the Public Network Settings for the Cisco Intercompany Media Engine Proxy, page 64-20.

Step 4 Specify the media termination address settings of the Cisco UMC. See Configuring the Media Termination Address for Unified Communication Proxies, page 64-10.

Step 5 Configure the local-side certificate management, namely the certificates that are exchanged between the local Cisco Unified Communications Manager servers and the security appliance. See Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 64-21.

Step 6 Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. This certificate is presented to remote servers so that they can authenticate the ASA as a trusted server. See Configuring the Remote-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 64-22.

The wizard completes by displaying a summary of the configuration created for the Cisco Intercompany Media Engine.

Configuring the Topology for the Cisco Intercompany Media Engine Proxy

Step 1 Select the topology of your ICME deployment by clicking one of the following options:

- All Internet traffic flows through the ASA radio button. This option is also referred to as a basic deployment.
- This ASA is off the path of the regular Internet traffic. This option is also referred to as an off-path deployment.

Step 2 Click **Next**.

Basic Deployment

In a basic deployment, the Cisco Intercompany Media Engine Proxy sits in-line with the Internet firewall such that all Internet traffic traverses the ASA. In this deployment, a single Cisco UCM or a Cisco UCM cluster is centrally deployed within the enterprise, along with a Cisco Intercompany Media Engine server (and perhaps a backup). A single Internet connection traverses the ASA, which is enabled with the Cisco Intercompany Media Engine Proxy.

The ASA sits on the edge of the enterprise and inspects SIP signaling by creating dynamic SIP trunks between enterprises.

Off-path Deployment

In an off path deployment, inbound and outbound Cisco Intercompany Media Engine calls pass through an ASA enabled with the Cisco Intercompany Media Engine Proxy. The ASA is located in the DMZ and configured to support primarily Cisco Intercompany Media Engine. Normal Internet facing traffic does not flow through this ASA.

For all inbound calls, the signaling is directed to the ASA because destined Cisco UCMs are configured with the global IP address on the ASA. For outbound calls, the called party could be any IP address on the Internet; therefore, the ASA is configured with a mapping service that dynamically provides an internal IP address on the ASA for each global IP address of the called party on the Internet.

Cisco UCM sends all outbound calls directly to the mapped internal IP address on the ASA instead of the global IP address of the called party on the Internet. The ASA then forwards the calls to the global IP address of the called party.

Note When you configure the Cisco Intercompany Media Engine for an off-path deployment, you must ensure that the public IP addresses and ports of the Cisco Unified Communications Manager servers and the public IP address for the media termination address are accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of the requirements.

Configuring the Private Network Settings for the Cisco Intercompany Media Engine Proxy

When configuring the Cisco Intercompany Media Engine Proxy, you specify settings to define the private network topology, such the private network interface, the IP addresses of the Cisco Unified Communications servers, and ticket verification. Additionally, when the Cisco Unified Communications servers are operating in secure mode, you specify the X.509 subject name for the Cisco Intercompany Media Engine Proxy,

The values that you specify in this page generate the following configuration settings for the Cisco Intercompany Media Engine Proxy:

- The list of Cisco Unified Communications servers
- The ticket epoch and password used by the Cisco Intercompany Media Engine Proxy
- For an off-path deployment only, the mapping service on the same interface as the Cisco Unified Communications server

Step 1 To configure the Cisco Intercompany Media Engine Proxy as part of a basic deployment, select the interface that connects to the local Cisco Unified Communications servers.

Or

To configure the Cisco Intercompany Media Engine Proxy as part of an off-path deployment, complete the following steps:

a. From the Listening Interface drop-down list, choose the interface on which the ASA listens for the mapping requests.

b. In the Port field, enter a number between 1024 and 65535 as the TCP port on which the ASA listens for the mapping requests. The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

c. From the UC-IME Interface drop-down list, choose the interface that the ASA uses to connect to the remote ASA that is enabled with the Cisco Intercompany Media Engine Proxy.

Note In a basic and an off-path deployment, all Cisco Unified Communications servers must be on the same interface.

Step 2 In the Unified CM Servers area, the wizard displays the private IP address, public IP address, and security mode of any Cisco Unified Communications server configured on the ASA. If necessary, click **Add** to add a Cisco Unified Communications server. You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled.

Step 3 In the Ticket Epoch field, enter a integer from 1-255.

The epoch indicates the number of times that password has changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password. Typically, you increment the epoch sequentially; however, the security appliance allows you to choose any value when you update the epoch.

If you change the epoch value, the current password is invalidated and you must enter a new password.

Step 4 In the Ticket Password field, enter a minimum of 10 and a maximum of 64 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character. The ticket password is stored onto flash.

Note We recommend a password of at least 20 characters. Only one password can be configured at a time.

The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Step 5 In the Confirm Password field, reenter the password.

Step 6 In the X.509 Subject Name field, enter the distinguished name (DN) of the local enterprise. The name that you enter must match the name configured for the Cisco Unified Communications servers in the cluster. See the Cisco Unified Communications server documentation for information.

Step 7 Click **Next**.

Adding a Cisco Unified Communications Manager Server for the UC-IME Proxy

You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine Proxy that has a SIP trunk enabled.

Step 1 Enter the private IP address and port number (in the range 5000-6000) for the Cisco UCM server.

Step 2 In the Address Translation area, enter the public IP address for the Cisco UCM server.

Step 3 If necessary, enter the port number for the public IP address by clicking the Translate address and port radio button and entering a number (in the range 5000-6000) in the Port field.

Step 4 In the Security Mode area, click the Secure or Non-secure radio button. Specifying secure for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS.

If you specify that some of the Cisco UCM servers are operating in secure mode, the Unified Communications Wizard includes a step in the proxy configuration to generate certificates for the local-side communication between the ASA and that Cisco UCM server. See Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 64-21.

Step 5 Click **OK**.

Configuring the Public Network Settings for the Cisco Intercompany Media Engine Proxy

The public network configuration depends on the deployment scenario you selected in the topology step of this wizard. Specifically, when you are configuring the UC-IME proxy as part of an off-path deployment, this step of the wizard displays fields for address translation, requiring that you specify the private IP address for the UC-IME proxy. Specifying this private IP address, translates IP addresses for inbound traffic.

In an off-path deployment, any existing ASA that you have deployed in your environment are not capable of transmitting Cisco Intercompany Media Engine traffic. Therefore, off-path signaling requires that outside addresses translate to an inside (private) IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address.

The values that you specify in this page generate the following configuration settings for the Cisco Intercompany Media Engine Proxy:

- Static PAT for the Cisco Unified Communications servers
- Access lists for traffic between the local and the remote servers

Step 1 In the Configure public network area, choose an interface from the Interface drop-down list.

Step 2 When configuring an off-path deployment, in the Address Translation area, specify whether to use the private IP address for the public network.

Or

Click the Specify IP address radio button and enter an IP address in the field.

Step 3 Click **Next**.

Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy

Completing this step of the wizard generates a self-signed certificate for the ASA. The server proxy certificate is automatically generated using the subject name provided in an earlier step of this wizard.

The wizard supports using self-signed certificates only.

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The certificates are used by the security appliance and the Cisco UCMs to authenticate each other, respectively, during TLS handshakes.

The ASA's identity certificate is exported, and then needs to be installed on each Cisco Unified Communications Manager (UCM) server in the cluster with the proxy and each identity certificate from the Cisco UCMs need to be installed on the security appliance.

This step in the Unified Communications Wizard only appears when the UC-IME proxy that you are creating has at least one secure Cisco Unified Communications Manager server defined. See Configuring the Topology for the Cisco Intercompany Media Engine Proxy, page 64-17 for information.

Step 1 In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog boxes appear indicating that the enrollment seceded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

Note

- If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.
- When using the wizard to configure the Cisco Intercompany Media Engine Proxy, the wizard only supports installing self-signed certificates.

Step 2 Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 64-23.

Step 3 In the Local Unified CM's Certificate area, click **Install Local Unified CM's Certificate**. The Install Certificate dialog appears.

Step 4 Locate the file containing the certificate from the Cisco Unified Communications Manager server or paste the certificate details in the dialog box. See Installing a Certificate, page 64-23. You must install the certificate from each Cisco Unified Communications Manager server in the cluster.

Step 5 Click **Next**.

Note See the Cisco Intercompany Media Engine server documentation for information on how to export the certificate for this server.

Configuring the Remote-Side Certificates for the Cisco Intercompany Media Engine Proxy

Establishing a trust relationship cross enterprises or across administrative domains is key. Cross enterprises you must use a trusted third-party CA (such as, VeriSign). The ASA obtains a certificate with the FQDN of the Cisco Unified Communications Manager server (certificate impersonation).

For the TLS handshake, the two entities could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. Both entities enroll with the CAs. The ASA as the TLS proxy must be trusted by both entities. The ASA is always associated with one of the enterprises. Within that enterprise, the entity and the ASA could authenticate each other via a local CA, or by using self-signed certificates.

To establish a trusted relationship between the ASA and the remote entity, the ASA can enroll with the CA on behalf of the local enterprise. In the enrollment request, the local Cisco UCM identity (domain name) is used.

To establish the trust relationship, the ASA enrolls with the third party CA by using the Cisco Unified Communications Manager server FQDN as if the security appliance is the Cisco UCM.

Note If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 3.

Step 1 In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 64-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 64-25.

Step 2 In the ASA's Identity Certificate area, click **Install ASA's Identity Certificate**. Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 64-26.

Step 3 In the Remote Server's CA's Certificate area, click **Install Remote Server's CA's Certificate**. Installing the root certificates of the CA for the remote servers is necessary so that the ASA can determine that the remote servers are trusted.

The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 64-23.

Note You must install the root certificates only when the root certificates for the remote servers are received from a CA other than the one that provided the identity certificate for the ASA

Step 4 Click **Next**.

The wizard completes by displaying a summary of the configuration created for the Cisco Intercompany Media Engine.

Working with Certificates in the Unified Communication Wizard

This section includes the following topics:

- Exporting an Identity Certificate, page 64-23
- Installing a Certificate, page 64-23
- Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 64-24
- Saving the Identity Certificate Request, page 64-25
- Installing the ASA Identity Certificate on the Mobility Advantage Server, page 64-26
- Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 64-26

Exporting an Identity Certificate

The Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy require that you export the ASA identity certificate to install on the Cisco Mobility Advantage server, Cisco Presence Federation server, and Cisco Unified Communications server, respectfully.

You use the wizard to export a self-signed identity certificate. The identity certificate has all associated keys and is in PKCS12 format, which is the public key cryptography standard. When configuring a Unified Communications proxy by using the wizard, you click the Generate and Export ASA's Identify Certificate button while in the local-side or server-side certificate management step of the wizard. The Export certificate dialog box appears.

From the Export certificate dialog box, perform these steps:

Step 1 Enter the name of the PKCS12 format file to use in exporting the certificate configuration. Alternatively, click Browse to display the Export ID Certificate File dialog box to find the file to which you want to export the certificate configuration.

Step 2 Click Export Certificate to export the certificate configuration.

An information dialog box appears informing you that the certificate configuration file has been successfully exported to the location that you specified.

You complete the configuration of the Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy, you must import the generated ASA identify certificate in to the Cisco Mobility Advantage server, Cisco Presence Federation server, and Cisco Unified Communications server, respectfully, depending on which proxy you are configuring.

See the documentation for the for each of these products for information about importing an identity certificate into each.

Installing a Certificate

When configuring certificates for the Phone Proxy, Cisco Mobility Advantage Proxy, the Cisco Presence Federation Proxy, and Cisco Intercompany Media Engine Proxy, you must install the certificates from the Cisco Unified Communications Manager servers, the Cisco Mobility Advantage server, the Cisco

Presence Federation server, and the Cisco Unified Communications Manager servers, respectively, on the ASA. See the documentation for each of these products for information about obtaining the identity certificates from each.

When configuring the Cisco Phone Proxy, if LSC provisioning is required or you have LSC enabled IP phones, you must install the CAPF certificate from the Cisco UCM on the ASA. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA. See Enabling Certificate Authority Proxy Function (CAPF) for IP Phones, page 64-8.

Additionally, when configuring the Cisco Mobility Advantage Proxy, you use the Install Certificate dialog box to install the root certificate received from the certificate authority. The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

Note When using the wizard to configure the Unified Communications proxies, the wizard only supports installing self-signed certificates.

From the Install Certificate dialog box, perform these steps:

Step 1 Perform one of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- To enroll manually, click the **Paste certificate in PEM format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2 Click **Install Certificate**.

An information dialog box appears informing you that the certificate was installed on the ASA successfully.

Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy

When configuring certificates for the Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy, you must generate and identity certificate request for the ASA.

Note If the ASA already has a signed identity certificate, you do not need to generate a CSR and can proceed directly to installing this certificate on the ASA. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 64-26 and Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 64-26 for the steps to install the identity certificate.

The identify certificate that you receive is presented to the following entities for each of the Unified Communication Proxies:

- Unified Mobile Communicator clients for the Cisco Mobility Advantage Proxy

- Remote Presence Federation servers for the Cisco Presence Federation Proxy
- The remote ASAfor the Cisco Intercompany Media Engine Proxy

Before generating the CSR, you can enter additional parameters.

When configuring a Unified Communications proxy by using the wizard, you click the Generate CSR button while in the client-side or remote-side certificate management step of the wizard. The CSR Parameters dialog box appears.

In the CSR Parameters dialog box, perform the following steps:

Step 1 From the Key Pair Size drop-down list, choose the size required for you certificate.

The key size that you select depends on the level of security that you want to configure and on any limitations imposed by the CA from which you are obtaining the certificate. The larger the number that you select, the higher the security level will be for the certificate. Most CAs recommend 2048 for the key modulus size; however, GoDaddy requires a key modulus size of 2048.

Step 2 (Cisco Intercompany Media Engine Proxy only) In the CN field, enter the domain name used by your enterprise or network. The subject DN you configure for the Cisco Intercompany Media Engine Proxy must match the domain name that set in the local Cisco Unified Communications Manager server.

Note For the Cisco Mobility Advantage Proxy and Cisco Presence Federation Proxy, the wizard provides the common name (CN), which is the FQDN of the Cisco Mobility Advantage server or Cisco Unified Presence server, respectively.

Step 3 In the Additional DN Attributes field, enter an attribute.

Or

Click **Select** to display the Additional DN Attributes dialog box.

a. In the Additional DN Attributes dialog box, choose an attribute from the drop-down list.
b. Enter a value for the attribute.
c. Click Add. The attribute appears in the list.
d. Click OK to return to the CSR Parameters dialog box.

The value you added appears in the Additional DN Attributes field in the CSR Parameters dialog box.

Step 4 Click **OK**.

Saving the Identity Certificate Request

After successfully generating the identity certificate request for one of the Unified Communications proxies, the Identity Certificate Request dialog box appears and prompts you to save the request.

Step 1 In the Save CSR to File field, enter the CSR file name and path; for example, c:\asa-csr.txt.

Step 2 Click **OK**. An information dialog box appears indicating the CSR was saved successfully.

Step 3 Click **OK** to close the dialog and return to the wizard.

Submit the CSR to the certificate authority (CA), for example, by pasting the CSR text into the CSR enrollment page on the CA website.

When the CA returns the signed identity certificate, rerun the Unified Communications Wizard. From the client-side or remote-side certificate management step of the wizard, click **Install ASA's Identity Certificate**. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 64-26 and Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 64-26 for the steps to install the identity certificate.

Installing the ASA Identity Certificate on the Mobility Advantage Server

When configuring certificates for the Cisco Mobility Advantage Proxy, you must install the ASA identity certificate on the Cisco Mobility Advantage server.

Typically, a certificate authority returns two certificates: your signed identity certificate and the certificate authority's certificate (referred to as the root certificate). However, some certificate authorities (for example, VeriSign) might also send you an intermediate certificate.

The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

If the certificate authority provided an intermediate certificate, you must enter the certificate text in the Intermediate Certificate (If Applicable) area of the Install ASA's Identity Certificate dialog box.

For the Cisco Mobility Advantage Proxy, you install the root certificate in another dialog box. See Installing a Certificate, page 64-23 for the steps to install the root certificate.

Step 1 In the Intermediate Certificate (If Applicable) area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2 In the ASA's Identity Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 3 Click **Install Certificate**.

Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers

When configuring certificates for the Cisco Presence Federation Proxy and Cisco Intercompany Media Engine Proxy, you must install the ASA identity certificate and the root certificate on the Cisco Presence Federation server and Cisco Intercompany Media Engine server, respectively.

Typically, a certificate authority returns two certificates: your signed identity certificate and the certificate authority's certificate (referred to as the root certificate). The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

Step 1 In the Root CA's Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2 In the ASA's Identity Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 3 Click **Install Certificate**.

Configuring the Cisco Phone Proxy

This chapter describes how to configure the adaptive security appliance for Cisco Phone Proxy feature.

This chapter includes the following sections:

- Information About the Cisco Phone Proxy, page 65-1
- Licensing Requirements for the Phone Proxy, page 65-4
- Prerequisites for the Phone Proxy, page 65-6
- Phone Proxy Guidelines and Limitations, page 65-12
- Configuring the Phone Proxy, page 65-14
- Feature History for the Phone Proxy, page 65-22

Information About the Cisco Phone Proxy

The Cisco Phone Proxy on the ASA bridges IP telephony between the corporate IP telephony network and the Internet in a secure manner by forcing data from remote phones on an untrusted network to be encrypted.

Phone Proxy Functionality

Telecommuters can connect their IP phones to the corporate IP telephony network over the Internet securely via the phone proxy without the need to connect over a VPN tunnel as illustrated by Figure 65-1.

Figure 65-1 Phone Proxy Secure Deployment

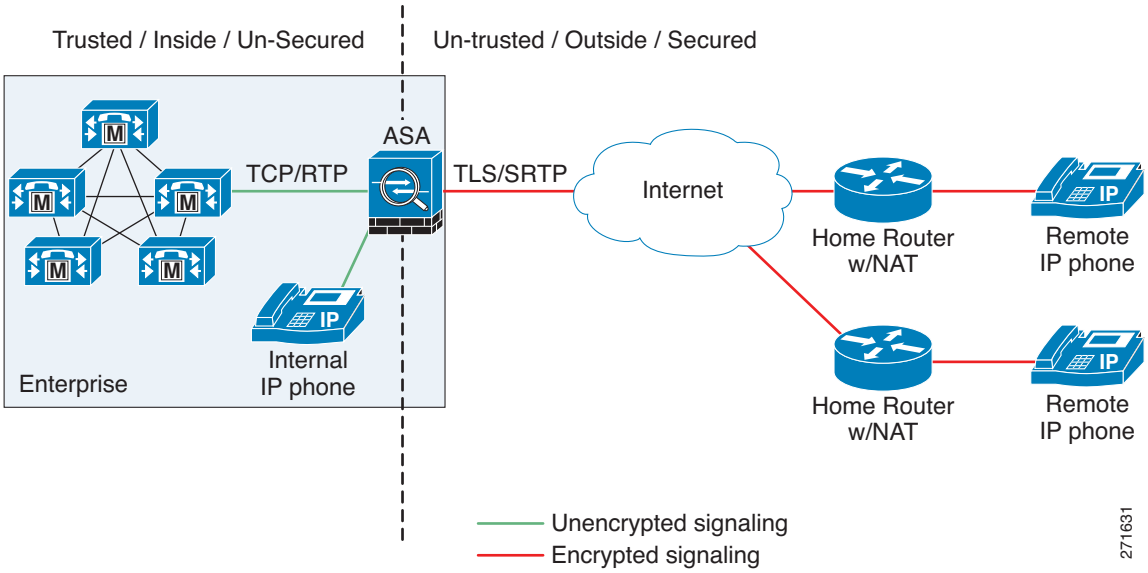

The phone proxy supports a Cisco UCM cluster in mixed mode or nonsecure mode. Regardless of the cluster mode, the remote phones that are capable of encryption are always forced to be in encrypted mode. TLS (signaling) and SRTP (media) are always terminated on the ASA. The ASA can also perform NAT, open pinholes for the media, and apply inspection policies for the SCCP and SIP protocols. In a nonsecure cluster mode or a mixed mode where the phones are configured as nonsecure, the phone proxy behaves in the following ways:

- The TLS connections from the phones are terminated on the ASA and a TCP connection is initiated to the Cisco UCM.
- SRTP sent from external IP phones to the internal network IP phone via the ASA is converted to RTP.

In a mixed mode cluster where the internal IP phones are configured as authenticated, the TLS connection is not converted to TCP to the Cisco UCM but the SRTP is converted to RTP.

In a mixed mode cluster where the internal IP phone is configured as encrypted, the TLS connection remains a TLS connection to the Cisco UCM and the SRTP from the remote phone remains SRTP to the internal IP phone.

Since the main purpose of the phone proxy is to make the phone behave securely while making calls to a nonsecure cluster, the phone proxy performs the following major functions:

- Creates the certificate trust list (CTL) file, which is used to perform certificate based authentication with remote phones.
- Modifies the IP phone configuration file when it is requested via TFTP, changes security fields from nonsecure to secure, and signs all files sent to the phone. These modifications secure remote phones by forcing the phones to perform encrypted signaling and media.
- Terminates TLS signaling from the phone and initiates TCP or TLS to Cisco UCM
- Inserts itself into the media path by modifying the Skinny and SIP signaling messages.
- Terminates SRTP and initiates RTP/SRTP to the called party.

Note As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via LSC provisioning. With LSC provisioning you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

Supported Cisco UCM and IP Phones for the Phone Proxy

Cisco Unified Communications Manager

The following release of the Cisco Unified Communications Manager are supported with the phone proxy:

- Cisco Unified CallManager Version 4.*x*
- Cisco Unified CallManager Version 5.0
- Cisco Unified CallManager Version 5.1
- Cisco Unified Communications Manager 6.1
- Cisco Unified Communications Manager 7.0

Cisco Unified IP Phones

The phone proxy supports these IP phone features:

- Enterprise features like conference calls on remote phones connected through the phone proxy
- XML services

The following IP phones in the Cisco Unified IP Phones 7900 Series are supported with the phone proxy:

- Cisco Unified IP Phone 7975
- Cisco Unified IP Phone 7971
- Cisco Unified IP Phone 7970
- Cisco Unified IP Phone 7965
- Cisco Unified IP Phone 7962
- Cisco Unified IP Phone 7961
- Cisco Unified IP Phone 7961G-GE
- Cisco Unified IP Phone 7960 (SCCP protocol support only)
- Cisco Unified IP Phone 7945
- Cisco Unified IP Phone 7942
- Cisco Unified IP Phone 7941
- Cisco Unified IP Phone 7941G-GE

- Cisco Unified IP Phone 7940 (SCCP protocol support only)
- Cisco Unified Wireless IP Phone 7921
- Cisco Unified Wireless IP Phone 7925

Note To support Cisco Unified Wireless IP Phone 7925, you must also configure MIC or LSC on the IP phone so that it properly works with the phone proxy.

- CIPC for softphones (CIPC versions with Authenticated mode only)

Note The Cisco IP Communicator is supported with the phone proxy VLAN Traversal in authenticated TLS mode. We do not recommend it for remote access because SRTP/TLS is not supported currently on the Cisco IP Communicator.

Note The ASA supports inspection of traffic from Cisco IP Phones running SCCP protocol version 19 and earlier.

Licensing Requirements for the Phone Proxy

The Cisco Phone Proxy feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5510	Base License and Security Plus License: 2 sessions. *Optional licenses: 24, 50, or 100 sessions.*
ASA 5520	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5540	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5550	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5580	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASA 5512-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*

Model	License Requirement[1]
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
   - Phone Proxy
   - Presence Federation Proxy
   - Encrypted Voice Inspection

   Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

   Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

   You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

   Note: For license part numbers ending in “K8” (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in “K9” (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

   Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

   You might also use SRTP encryption sessions for your connections:
   - For K8 licenses, SRTP sessions are limited to 250.
   - For K9 licenses, there is not limit.

   Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.
2. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

For more information about licensing, see Chapter 5, “Managing Feature Licenses.”

Prerequisites for the Phone Proxy

This section contains the following topics:

- Media Termination Instance Prerequisites, page 65-6
- Certificates from the Cisco UCM, page 65-7
- DNS Lookup Prerequisites, page 65-7
- Cisco Unified Communications Manager Prerequisites, page 65-7
- Access List Rules, page 65-7
- NAT and PAT Prerequisites, page 65-8
- Prerequisites for IP Phones on Multiple Interfaces, page 65-8
- 7960 and 7940 IP Phones Support, page 65-9
- Cisco IP Communicator Prerequisites, page 65-10
- Prerequisites for Rate Limiting TFTP Requests, page 65-10
- End-User Phone Provisioning, page 65-11

Media Termination Instance Prerequisites

The ASA must have a media termination instance that meets the following criteria:

- You must configure one media termination for each phone proxy on the ASA. Multiple media termination instances on the ASA are not supported.
- For the media termination instance, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.
- If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the ASA uses when communicating with IP phones.

  For example, if you had three interfaces on the ASA (one internal interface and two external interfaces) and only one of the external interfaces were used to communicate with IP phones, you would configure two media termination addresses: one on the internal interface and one on the external interface that communicated with the IP phones.
- Only one media-termination address can be configured per interface.
- The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.
- The IP address on an interface cannot be the same address as that interface on the ASA.
- The IP addresses cannot overlap with existing static NAT pools or NAT rules.
- The IP addresses cannot be the same as the Cisco UCM or TFTP server IP address.
- For IP phones behind a router or gateway, you must also meet this prerequisite. On the router or gateway, add routes to the media termination address on the ASA interface that the IP phones communicate with so that the phone can reach the media termination address.

Certificates from the Cisco UCM

Import the following certificates which are stored on the Cisco UCM. These certificates are required by the ASA for the phone proxy.

- Cisco_Manufacturing_CA
- CAP-RTP-001
- CAP-RTP-002
- CAPF certificate (Optional)

If LSC provisioning is required or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA.

Note You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

For example, the CA Manufacturer certificate is required by the phone proxy to validate the IP phone certificate.

DNS Lookup Prerequisites

- If you have an fully qualified domain name (FQDN) configured for the Cisco UCM rather than an IP address, you must configure and enable DNS lookup on the ASA.
- After configuring the DNS lookup, make sure that the ASA can ping the Cisco UCM with the configured FQDN.
- You must configure DNS lookup when you have a CAPF service enabled and the Cisco UCM is not running on the Publisher but the Publisher is configured with a FQDN instead of an IP address.

Cisco Unified Communications Manager Prerequisites

- The TFTP server must reside on the same interface as the Cisco UCM.
- The Cisco UCM can be on a private network on the inside but you need to have a static mapping for the Cisco UCM on the ASA to a public routable address.
- If NAT is required for Cisco UCM, it must be configured on the ASA, not on the existing firewall.

Access List Rules

If the phone proxy is deployed behind an existing firewall, access-list rules to permit signaling, TFTP requests, and media traffic to the phone proxy must be configured.

If NAT is configured for the TFTP server or Cisco UCMs, the translated “global” address must be used in the access lists.

Table 65-1 lists the ports that are required to be configured on the existing firewall:

Table 65-1 Port Configuration Requirements

Address	Port	Protocol	Description
Media Termination	1024-65535	UDP	Allow incoming SRTP
TFTP Server	69	UDP	Allow incoming TFTP
Cisco UCM	2443	TCP	Allow incoming secure SCCP
Cisco UCM	5061	TCP	Allow incoming secure SIP
CAPF Service (on Cisco UCM)	3804	TCP	Allow CAPF service for LSC provisioning

Note All these ports are configurable on the Cisco UCM, except for TFTP. These are the default values and should be modified if they are modified on the Cisco UCM. For example, 3804 is the default port for the CAPF Service. This default value should be modified if it is modified on the Cisco UCM.

NAT and PAT Prerequisites

NAT Prerequisites

- If NAT is configured for the TFTP server, the NAT configuration must be configured prior to configuring the TFTP Server for the phone proxy.
- If NAT is configured for the TFTP server or Cisco UCMs, the translated “global” address must be used in the access lists.

PAT Prerequisites

- When the Skinny inspection global port is configured to use a non-default port, then you must configure the nonsecure port as the `global_sccp_port+443`.

  Therefore, if *global_sccp_port* is 7000, then the global secure SCCP port is 7443. Reconfiguring the port might be necessary when the phone proxy deployment has more than one Cisco UCM and they must share the interface IP address or a global IP address.

Note Both PAT configurations—for the nonsecure and secure ports—must be configured.

- When the IP phones must contact the CAPF on the Cisco UCM and the Cisco UCM is configured with static PAT (LCS provisioning is required), you must configure static PAT for the default CAPF port 3804.

Prerequisites for IP Phones on Multiple Interfaces

When IP phones reside on multiple interfaces, the phone proxy configuration must have the correct IP address set for the Cisco UCM in the CTL file.

See the following example topology for information about how to correctly set the IP address:

```
phones --- (dmz)-----|
                     |----- ASA PP --- (outside Internet) --- phones
phones --- (inside)--|
```

In this example topology, the following IP address are set:

- Cisco UCM on the inside interface is set to 10.0.0.5
- The DMZ network is 192.168.1.0/24
- The inside network is 10.0.0.0/24

The Cisco UCM is mapped with different global IP addresses from DMZ > outside and inside interfaces > outside interface.

In the CTL file, the Cisco UCM must have two entries because of the two different IP addresses. For example, if the static statements for the Cisco UCM are as follows:

```
object network obj-10.0.0.5-01
    host 10.0.0.5
    nat (inside,outside) static 209.165.202.129
object network obj-10.0.0.5-02
    host 10.0.0.5
    nat (inside,dmz) static 198.168.1.2
```

There must be two CTL file record entries for the Cisco UCM:

```
record-entry cucm trustpoint cucm_in_to_out address 209.165.202.129
record-entry cucm trustpoint cucm_in_to_dmz address 192.168.1.2
```

7960 and 7940 IP Phones Support

- An LSC must be installed on these IP phones because they do not come pre installed with a MIC. Install the LSC on each phone before using them with the phone proxy to avoid opening the nonsecure SCCP port for the IP phones to register in nonsecure mode with the Cisco UCM.

  See the following document for the steps to install an LSC on IP phones:

  http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/security/7_0_1/secugd/secucapf.html#wp1093518

Note If an IP phone already has an LSC installed on it from a different Cisco UCM cluster, delete the LSC from the different cluster and install an LSC from the current Cisco UCM cluster.

Note You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

- The CAPF certificate must be imported onto the ASA.
- The CTL file created on the ASA must be created with a CAPF record-entry.
- The phone must be configured to use only the SCCP protocol because the SIP protocol does not support encryption on these IP phones.
- If LSC provisioning is done via the phone proxy, you must add an ACL to allow the IP phones to register with the Cisco UCM on the nonsecure port 2000.

Cisco IP Communicator Prerequisites

To configure Cisco IP Communicator (CIPC) with the phone proxy, you must meet the following prerequisites:

- Go to Configuration > Firewall > Unified Communications > Phone Proxy and select the "Enable CICP security mode authentication" check box under the Call Manager and Phone Settings area.
- Create an ACL to allow CIPC to register with the Cisco UCM in nonsecure mode.
- Configure null-sha1 as one of the SSL encryption ciphers.

Current versions of Cisco IP Communicator (CIPC) support authenticated mode and perform TLS signaling but not voice encryption.

Because CIPC requires an LSC to perform the TLS handshake, CIPC needs to register with the Cisco UCM in nonsecure mode using cleartext signaling. To allow the CIPC to register, create an ACL that allows the CIPC to connect to the Cisco UCM on the nonsecure SIP/SCCP signalling ports (5060/2000).

Note You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

CIPC uses a different cipher when doing the TLS handshake and requires the null-sha1 cipher and SSL encryption be configured. To add the null-shal cipher, use the show run all ssl command to see the output for the ssl encryption command and add null-shal to the end of the SSL encryption list.

Note When used with CIPC, the phone proxy does not support end-users resetting their device name in CIPC (Preferences > Network tab > Use this Device Name field) or Administrators resetting the device name in Cisco Unified CM Administration console (Device menu > Phone Configuration > Device Name field). To function with the phone proxy, the CIPC configuration file must be in the format: SEP<mac_address>.cnf.xml. If the device name does not follow this format (SEP<mac_address>), CIPC cannot retrieve its configuration file from Cisco UMC via the phone proxy and CIPC will not function.

Prerequisites for Rate Limiting TFTP Requests

In a remote access scenario, we recommend that you configure rate limiting of TFTP requests because any IP phone connecting through the Internet is allowed to send TFTP requests to the TFTP server.

To configure rate limiting of TFTP requests, configure the **police** command in the Modular Policy Framework. See the command reference for information about using the **police** command.

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow can take over the entire resource. When traffic exceeds the maximum rate, the ASA drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Rate Limiting Configuration Example

The following example describes how you configure rate limiting for TFTP requests by using the **police** command and the Modular Policy Framework.

Begin by determining the conformance rate that is required for the phone proxy. To determine the conformance rate, use the following formula:

```
X * Y * 8
```

Where

X = requests per second

Y = size of each packet, which includes the L2, L3, and L4 plus the payload

Therefore, if a rate of 300 TFTP requests/second is required, then the conformance rate would be calculated as follows:

```
300 requests/second * 80 bytes * 8 = 192000
```

To control which hosts can ping the media termination address, create an ICMP rule. Go to Configuration > Device Management > Management Access > ICMP and click the Add button.

End-User Phone Provisioning

The phone proxy is a transparent proxy with respect to the TFTP and signaling transactions. If NAT is not configured for the Cisco UCM TFTP server, then the IP phones need to be configured with the Cisco UCM cluster TFTP server address.

If NAT is configured for the Cisco UCM TFTP server, then the Cisco UCM TFTP server global address is configured as the TFTP server address on the IP phones.

Ways to Deploy IP Phones to End Users

In both options, deploying a remote IP phone behind a commercial Cable/DSL router with NAT capabilities is supported.

Option 1 (Recommended)

Stage the IP phones at corporate headquarters before sending them to the end users:

- The phones register inside the network. IT ensures there are no issues with the phone configurations, image downloads, and registration.
- If Cisco UCM cluster was in mixed mode, the CTL file should be erased before sending the phone to the end user.

Advantages of this option are:

- Easier to troubleshoot and isolate problems with the network or phone proxy because you know whether the phone is registered and working with the Cisco UCM.
- Better user experience because the phone does not have to download firmware from over a broadband connection, which can be slow and require the user to wait for a longer time.

Option 2

Send the IP phone to the end user. When using option 2, the user must be provided instructions to change the settings on phones with the appropriate Cisco UCM and TFTP server IP address.

Note As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via LSC provisioning. With LSC provisioning you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register

in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

Phone Proxy Guidelines and Limitations

This section includes the following topics:

- General Guidelines and Limitations, page 65-12
- Media Termination Address Guidelines and Limitations, page 65-13

General Guidelines and Limitations

The phone proxy has the following general limitations:

- Only one phone proxy instance can be configured on the ASA by using the **phone-proxy** command. See the command reference for information about the **phone-proxy** command. See also Creating the Phone Proxy Instance, page 65-18.
- The phone proxy only supports one Cisco UCM cluster. See Creating the CTL File, page 65-14 for the steps to configure the Cisco UCM cluster for the phone proxy.
- The phone proxy is not supported when the ASA is running in transparent mode or multiple context mode.
- When a remote IP phone calls an invalid internal or external extension, the phone proxy does not support playing the annunciator message from the Cisco UCM. Instead, the remote IP phone plays a fast busy signal instead of the annunciator message "Your call cannot be completed ..." However, when an internal IP phone dials in invalid extension, the annunciator messages plays "Your call cannot be completed ..."
- Packets from phones connecting to the phone proxy over a VPN tunnel are not inspected by the ASA inspection engines.
- The phone proxy does not support IP phones sending Real-Time Control Protocol (RTCP) packets through the ASA. Disable RTCP packets in the Cisco Unified CM Administration console from the Phone Configuration page. See your Cisco Unified Communications Manager (CallManager) documentation for information about setting this configuration option.
- When used with CIPC, the phone proxy does not support end-users resetting their device name in CIPC (Preferences > Network tab > Use this Device Name field) or Administrators resetting the device name in Cisco Unified CM Administration console (Device menu > Phone Configuration > Device Name field). To function with the phone proxy, the CIPC configuration file must be in the format: SEP<mac_address>.cnf.xml. If the device name does not follow this format (SEP<mac_address>), CIPC cannot retrieve its configuration file from Cisco UMC via the phone proxy and CIPC will not function.
- The phone proxy does not support IP phones sending SCCP video messages using Cisco VT Advantage because SCCP video messages do not support SRTP keys.
- For mixed-mode clusters, the phone proxy does not support the Cisco Unified Call Manager using TFTP to send encrypted configuration files to IP phones through the ASA.

- Multiple IP phones behind one NAT device must be configured to use the same security mode.

  When the phone proxy is configured for a mixed-mode cluster and multiple IP phones are behind one NAT device and registering through the phone proxy, all the SIP and SCCP IP phones must be configured as authenticated or encrypted, or all as non-secure on the Unified Call Manager.

  For example, if there are four IP phones behind one NAT device where two IP phones are configured using SIP and two IP phones are configured using SCCP, the following configurations on the Unified Call Manager are acceptable:

  – Two SIP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, or both in encrypted mode

    Two SCCP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, or both in encrypted mode

  – Two SIP IP phones: both in non-secure mode

    Two SCCP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, both in encrypted mode

  – Two SIP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, both in encrypted mode

    Two SCCP IP phones: both in non-secure mode

  This limitation results from the way the application-redirect rules (rules that convert TLS to TCP) are created for the IP phones.

Media Termination Address Guidelines and Limitations

The phone proxy has the following limitations relating to configuring the media-termination address:

- When configuring the media-termination address, the phone proxy does not support having internal IP phones (IP phones on the inside network) being on a different network interface from the Cisco UCM unless the IP phones are forced to use the non-secure Security mode.

  When internal IP phones are on a different network interface than the Cisco UCM, the IP phones signalling sessions still go through ASA; however, the IP phone traffic does not go through the phone proxy. Therefore, Cisco recommends that you deploy internal IP phones on the same network interface as the Cisco UMC.

  If the Cisco UMC and the internal IP phones must be on different network interfaces, you must add routes for the internal IP phones to access the network interface of the media-termination address where Cisco UMC resides.

  When the phone proxy is configured to use a global media-termination address, all IP phones see the same global address, which is a public routable address.

- If you decide to configure a media-termination address on interfaces (rather than using a global interface), you must configure a media-termination address on at least two interfaces (the inside and an outside interface) before applying the phone-proxy service policy. Otherwise, you will receive an error message when enabling the Phone Proxy with SIP and Skinny Inspection.

- The phone proxy can use only one type of media termination instance at a time; for example, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.

Configuring the Phone Proxy

This section includes the following topics:

- Task Flow for Configuring the Phone Proxy, page 65-14
- Creating the CTL File, page 65-14
- Adding or Editing a Record Entry in a CTL File, page 65-16
- Creating the Media Termination Instance, page 65-17
- Creating the Phone Proxy Instance, page 65-18
- Adding or Editing the TFTP Server for a Phone Proxy, page 65-20
- Configuring Linksys Routers with UDP Port Forwarding for the Phone Proxy, page 65-21

Task Flow for Configuring the Phone Proxy

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Configuring the Phone Proxy requires the following steps:

Step 1: Create the CTL file. See Creating the CTL File, page 65-14.

Step 2: Create the TLS Proxy instance to handle the encrypted signaling. See Adding a TLS Proxy Instance, page 66-9.

Step 3: Create the Phone Proxy instance. See the "Creating the Phone Proxy Instance" section on page 65-18.

Step 4: Configure the media termination address for the Phone Proxy. See Creating the Media Termination Instance, page 65-17.

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 59-20 and Skinny (SCCP) Inspection, page 59-31.

Creating the CTL File

Create a Certificate Trust List (CTL) file that is required by the Phone Proxy. Specify the certificates needed by creating a new CTL file or by specifying the path of an exiting CTL file to parse from Flash memory.

Create trustpoints and generate certificates for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. The certificates are used in creating the CTL file. You need to create trustpoints for each CUCM (primary and secondary if a secondary CUCM is used) and TFTP server in the network. The trustpoints need to be in the CTL file for the phones to trust the CUCM.

Create the CTL File that will be presented to the IP phones during the TFTP. The address must be the translated or global address of the TFTP server or CUCM if NAT is configured.

When the file is created, it creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP**_*ctl-instance_filename*.

Note When a CTL file instance is assigned to the Phone Proxy, you cannot modify it in the CTL File pane and the pane is disabled. To modify a CTL File that is assigned to the Phone Proxy, go to the Phone Proxy pane (Configuration > Firewall > Unified Communications > Phone Proxy), and deselect the Use the Certificate Trust List File generated by the CTL instance check box.

Use the Create a Certificate Trust List (CTL) File pane to create a CTL file for the Phone Proxy. This pane creates the CTL file that is presented to the IP phones during the TFTP handshake with the ASA. For a detailed overview of the CTL file used by the Phone Proxy, see the "Creating the CTL File" section on page 65-14.

The Create a Certificate Trust List (CTL) File pane is used to configure the attributes for generating the CTL file. The name of the CTL file instance is generated by the ASDM. When the user tries to edit the CTL file instance configuration, the ASDM automatically generates the **shutdown** CLI command first and the **no shutdown** CLI command as the last command.

This pane is available from the Configuration > Firewall > Unified Communications > CTL File pane.

Step 1 Open the Configuration > Firewall > Unified Communications > CTL File pane.

Step 2 Check the Enable Certificate Trust List File check box to enable the feature.

Step 3 To specify the CTL file to use for the Phone Proxy, perform one of the following:

- If there is an existing CTL file available, download the CTL file to Flash memory by using the File Management Tool in the ASDM Tools menu. Select the Use certificates present in the CTL stored in flash radio button and specify the CTL file name and path in the text box.

  Use an existing CTL file to install the trustpoints for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. If you have an existing CTL file that contains the correct IP addresses of the entities (namely, the IP address that the IP phones use for the CUCM or TFTP servers), you can be use it to create a new CTL file. Store a copy of the existing CTL file to Flash memory and rename it something other than `CTLFile.tlv`

- If there is no existing CTL file available, select Create new CTL file radio button.

  Add Record entries for each entity in the network such as CUCM, TFTP, and CUCM-TFTP option by clicking **Add**. The Add Record Entry dialog box opens. See Adding or Editing a Record Entry in a CTL File, page 65-16.

Step 4 Specify the number SAST certificate tokens required. The default is 2. maximum allowed is 5.

Because the Phone Proxy generates the CTL file, it needs to create the System Administrator Security Token (SAST) key to sign the CTL file itself. This key can be generated on the ASA. A SAST is created as a self-signed certificate. Typically, a CTL file contains more than one SAST. In case a SAST is not recoverable, the other one can be used to sign the file later.

Step 5 Click **Apply** to save the CTL file configuration settings.

Adding or Editing a Record Entry in a CTL File

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add/Edit Record Entry dialog box to specify the trustpoints to be used for the creation of the CTL file.

Note You can edit an entry in the CTL file by using the Edit Record Entry dialog box; however, changing a setting in this dialog box does not change related settings for the phone proxy. For example, editing the IP address for the CUCM or TFTP servers in this dialog changes the setting only in the CTL file and does not change the actual addresses of those servers or update the address translations required by the phone proxy.
To modify CTL file settings, we strongly recommend you re-run the Unified Communications Wizard to edit CTL file settings and ensure proper synchronization with all phone proxy settings.

Add additional record-entry configurations for each entity that is required in the CTL file.

Step 1 Open the Configuration > Firewall > Unified Communications > CTL File pane.

Step 2 Check the Enable Certificate Trust List File check box to enable the feature.

Step 3 In the Type field, specify the type of trustpoint to create:

- cucm: Specifies the role of this trustpoint to be CCM. Multiple CCM trustpoints can be configured.
- cucm-tftp: Specifies the role of this trustpoint to be CCM+TFTP. Multiple CCM+TFTP trustpoints can be configured.
- tftp: Specifies the role of this trustpoint to be TFTP. Multiple TFTP trustpoints can be configured.
- capf: Specifies the role of this trustpoint to be CAPF. Only one CAPF trustpoint can be configured.

Step 4 In the Host field, specify the IP address of the trustpoint. The IP address you specify must be the global address of the TFTP server or CUCM if NAT is configured. The global IP address is the IP address as seen by the IP phones because it will be the IP address used for the CTL record for the trustpoint.

Step 5 In the Certificate field, specify the Identity Certificate for the record entry in the CTL file. You can create a new Identity Certificate by clicking **Manage**. The Manage Identify Certificates dialog box opens. See the "Configuring Identity Certificates Authentication" section on page 49-22.

You can add an Identity Certificate by generating a self-signed certificate, obtaining the certificate through SCEP enrollment, or by importing a certificate in PKCS-12 format. Choose the best option based on the requirements for configuring the CTL file.

Step 6 (Optional) In the Domain Name field, specify the domain name of the trustpoint used to create the DNS field for the trustpoint. This is appended to the Common Name field of the Subject DN to create the DNS Name. The domain name should be configured when the FQDN is not configured for the trustpoint. Only one domain-name can be specified.

Note If you are using domain names for your CUCM and TFTP server, you must configure DNS lookup on the ASA. Add an entry for each of the outside interfaces on the ASA into your DNS server, if such entries are not already present. Each ASA outside IP address should have a DNS entry associated with it for

lookups. These DNS entries must also be enabled for Reverse Lookup. Additionally, define your DNS server IP address on the ASA; for example: `dns name-server 10.2.3.4` (IP address of your DNS server).

Creating the Media Termination Instance

Create the media termination instance that you will use in the phone proxy.

The media termination address you configure must meet the requirements as described in Media Termination Instance Prerequisites, page 65-6.

Note In versions before 8.2(1), you configured one media-termination address (MTA) on the outside interface of the adaptive security appliance where the remote Cisco IP phones were located. In Version 8.2(1) and later, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.

As a result of this enhancement, the old configuration has been deprecated. You can continue to use the old configuration if desired. However, if you need to change the configuration at all, only the new configuration method is accepted; you cannot later restore the old configuration. If you need to maintain downgrade compatibility, you should keep the old configuration as is.

Step 1 Open the Configuration > Firewall > Unified Communications > Media Termination Address pane.

Step 2 Check the Enable Media Termination Address check box to enable the feature.

Step 3 In the Media Termination Address Settings area, specify whether to configure a media-termination address (MTA) per interface or to configure a global MTA. You can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.

- To configure an MTA per interface, click the Configure MTA per Interface radio button and click the **Add** button. In the dialog box that appears, specify the interface name and enter an IP address or hostname.

  If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the ASA uses when communicating with IP phones. The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.

  See Media Termination Instance Prerequisites, page 65-6 for the complete list of requirements that you must follow when creating the media termination instance and configuring the media termination addresses.

- To configure a global MTA, click the Configure global MTA on interface radio button and enter the IP address in the text box. See Media Termination Instance Prerequisites, page 65-6 for the complete list of requirements that you must follow when configuring a global media termination address.

Step 4 Specify the minimum and maximum values for the RTP port range for the media termination instance. The minimum port and the maximum port can be a value from 1024 to 65535.

Step 5 Click **Apply** to save the media termination address configuration settings.

Creating the Phone Proxy Instance

Create the phone proxy instance. To have a fully functional phone proxy, you must also complete additional tasks, such as creating the MTA and enabling SIP and SCCP (Skinny) inspection. See Task Flow for Configuring the Phone Proxy, page 65-14 for the complete list of tasks.

Prerequisites

You must have already created the CTL file and TLS proxy instance for the phone proxy.

See Creating the CTL File, page 65-14 and Adding a TLS Proxy Instance, page 66-9.

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Configure Phone Proxy pane to add a Phone Proxy.

This pane is available from the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 1 Open the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 2 Check the Enable Phone Proxy check box to enable the feature.

Step 3 Check the Apply MTA instance to Phone Proxy check box to add the media termination address to the Phone Proxy instance. You must have a media termination address instance configured. The configured address is added to the Phone Proxy instance.

Note To configure the media termination address, click the Configure MTA button. The Media Termination Address dialog box appears. Once you click the Add MTA instance to Phone Proxy check box, the media termination address instance cannot be modified and the button changes to View MTA Configuration. To change the media termination address, uncheck the Add MTA instance to Phone Proxy check box.

Step 4 If necessary, add a TFTP server for the Phone Proxy. To add a new TFTP server for the Phone Proxy, click **Add**. The Add TFTP Server dialog box opens. See Adding or Editing the TFTP Server for a Phone Proxy, page 65-20.

Note The TFTP server must reside on the same interface as the Cisco Unified Call Manager. Additionally, If NAT is configured for the TFTP server, the NAT configuration must be configured prior to configuring the specifying the TFTP server while creating the Phone Proxy instance.

Step 5 Specify the CTL File to use for the Phone Proxy by doing one of the following:

- To use an existing CTL File, check the Use the Certificate Trust List File generated by the CTL instance check box.
- To create a new CTL file for the Phone Proxy, click the link Generate Certificate Trust List File. The Create a Certificate Trust List (CTL) File pane opens. See "Creating the CTL File" section on page 65-14.

Step 6 To specify the security mode of the CUCM cluster, click one of the following options in the CUCM Cluster Mode field:

- Non-secure—Specifies the cluster mode to be in nonsecure mode when configuring the Phone Proxy feature.
- Mixed—Specifies the cluster mode to be in mixed mode when configuring the Phone Proxy feature.

Step 7 To configure the idle timeout after which the secure-phone entry is removed from the Phone Proxy database (the default is 5 minutes), enter a value in the format *hh*:*mm*:*ss*.

Since secure phones always request a CTL file upon bootup, the Phone Proxy creates a database that marks the phone as secure. The entries in the secure phone database are removed after a specified configured timeout. The entry timestamp is updated for each registration refresh the Phone Proxy receives for SIP phones and KeepAlives for SCCP phones.

Specify a value that is greater than the maximum timeout value for SCCP KeepAlives and SIP Register refresh. For example, if the SCCP KeepAlives are configured for 1 minute intervals and the SIP Register Refresh is configured for 3 minutes, configure this timeout value greater than 3 minutes.

Step 8 To preserve Call Manager configuration on the IP phones, check the Preserve the Call Manager's configuration on the phone... check box. When this check box is uncheck, the following service settings are disabled on the IP phones:

- PC Port
- Gratuitous ARP
- Voice VLAN access
- Web Access
- Span to PC Port

Step 9 To force Cisco IP Communicator (CIPC) softphones to operate in authenticated mode when CIPC softphones are deployed in a voice and data VLAN scenario, check the Enable CIPC security mode authentication check box.

Because CIPC requires an LSC to perform the TLS handshake, CIPC needs to register with the CUCM in nonsecure mode using cleartext signaling. To allow the CIPC to register, create an ACL that allows the CIPC to connect to the CUCM on the nonsecure SIP/SCCP signalling ports (5060/2000).

CIPC uses a different cipher when doing the TLS handshake and requires the null-sha1 cipher and SSL encryption be configured. To add the null-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 10 To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the <proxyServerURL> tag, do the following:

a. Check the Configure a http-proxy which would be written into the phone's config file... check box.

b. In the IP Address field, type the IP address of the HTTP proxy and the listening port of the HTTP proxy.

The IP address you enter should be the global IP address based on where the IP phone and HTTP proxy server is located. You can enter a hostname in the IP Address field when that hostname can be resolved to an IP address by the ASA (for example, DNS lookup is configured) because the ASA will resolve the hostname to an IP address. If a port is not specified, the default will be 8080.

c. In the Interface field, select the interface on which the HTTP proxy resides on the ASA.

Setting the proxy server configuration option for the Phone Proxy allows for an HTTP proxy on the DMZ or external network in which all the IP phone URLs are directed to the proxy server for services on the phones. This setting accommodates nonsecure HTTP traffic, which is not allowed back into the corporate network.

Step 11 Click **Apply** to save the Phone Proxy configuration settings.

Note After creating the Phone Proxy instance, you enable it with SIP and Skinny inspection. See SIP Inspection, page 59-20 and Skinny (SCCP) Inspection, page 59-31.

However, before you enable SIP and Skinny inspection for the Phone Proxy (which is done by applying the Phone Proxy to a service policy rule), the Phone Proxy must have an MTA instance, TLS Proxy, and CTL file assigned to it before the Phone Proxy can be applied to a service policy. Additionally, once a Phone Proxy is applied to a service policy rule, the Phone Proxy cannot be changed or removed.

Adding or Editing the TFTP Server for a Phone Proxy

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Note You can edit the TFTP server setting by using the Edit TFTP Server dialog box; however, changing a setting in this dialog box does not change related settings for the phone proxy. For example, editing the IP address for the TFTP server in this dialog does not change the setting in the CTL file and does not update the address translations required by the phone proxy.
To modify TFTP server settings, we strongly recommend you re-run the Unified Communications Wizard to ensure proper synchronization with all phone proxy settings.

Step 1 Open the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 2 Check the Enable Phone Proxy check box to enable the feature.

Step 3 To add or edit the TFTP Server information for the phone proxy, click the **Add** or **Edit** button. The Add/Edit TFTP Server dialog box appears.

Use the Add/Edit TFTP Server dialog box to specify the IP address of the TFTP server and the interface on which the TFTP server resides.

The Phone Proxy must have at least one CUCM TFTP server configured. Up to five TFTP servers can be configured for the Phone Proxy.

The TFTP server is assumed to be behind the firewall on the trusted network; therefore, the Phone Proxy intercepts the requests between the IP phones and TFTP server.

Note If NAT is configured for the TFTP server, the NAT configuration must be configured prior to specifying the TFTP server while creating the Phone Proxy instance.

Step 4 In the TFTP Server IP Address field, specify the address of the TFTP server. Create the TFTP server using the actual internal IP address.

Step 5 (Optional) In the Port field, specify the port the TFTP server is listening in on for the TFTP requests. This should be configured if it is not the default TFTP port 69.

Step 6 In the Interface field, specify the interface on which the TFTP server resides. The TFTP server must reside on the same interface as the Cisco Unified Call Manager (CUCM).

Step 7 Click OK to apply the settings.

Configuring Linksys Routers with UDP Port Forwarding for the Phone Proxy

When IP phones are behind a NAT-capable router, the router can be configured to forward the UDP ports to the IP address of the IP phone. Specifically, configure the router for UDP port forwarding when an IP phone is failing during TFTP requests and the failure is due to the router dropping incoming TFTP data packets. Configure the router to enable UDP port forwarding on port 69 to the IP phone.

As an alternative of explicit UDP forwarding, some Cable/DSL routers require you to designate the IP phone as a DMZ host. For Cable/DSL routers, this host is a special host that receives all incoming connections from the public network.

When configuring the phone proxy, there is no functional difference between an IP phone that has UDP ports explicitly forwarded or an IP phone designated as a DMZ host. The choice is entirely dependent upon the capabilities and preference of the end user.

Configuring Your Router

Your firewall/router needs to be configured to forward a range of UDP ports to the IP phone. This will allow the IP phone to receive audio when you make/receive calls.

Note Different Cable/DSL routers have different procedures for this configuration. Furthermore most NAT-capable routers will only allow a given port range to be forwarded to a single IP address

The configuration of each brand/model of firewall/router is different, but the task is the same. For specific instructions for your brand and model of router, please contact the manufacturer's website.

Linksys Routers

Step 1 From your web browser, connect to the router administrative web page. For Linksys, this is typically something like `http://192.168.1.1`.

Step 2 Click Applications & Gaming or the Port Forwarding tab (whichever is present on your router).

Step 3 Locate the table containing the port forwarding data and add an entry containing the following values:

Table 65-2 Port Forwarding Values to Add to Router

Application	Start	End	Protocol	IP Address	Enabled
IP phone	1024	65535	UDP	*Phone IP address*	**Checked**
TFTP	69	69	UDP	*Phone IP address*	**Checked**

Step 4 Click Save Settings. Port forwarding is configured.

Feature History for the Phone Proxy

Table 65-3 lists the release history for this feature.

Table 65-3 Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Phone Proxy	8.0(4)	The phone proxy feature was introduced. The Phone Proxy feature was accessible in ASDM by choosing the following options: Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Phone Proxy pane
NAT for the media termination address	8.1(2)	The Media Termination fields were removed from the Phone Proxy pane and added to the Media Termination pane: Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Media Termination Address pane

Configuring the TLS Proxy for Encrypted Voice Inspection

This chapter describes how to configure the adaptive security appliance for the TLS Proxy for Encrypted Voice Inspection feature.

This chapter includes the following sections:

- Information about the TLS Proxy for Encrypted Voice Inspection, page 66-1
- Licensing for the TLS Proxy, page 66-4
- Prerequisites for the TLS Proxy for Encrypted Voice Inspection, page 66-6
- Configuring the TLS Proxy for Encrypted Voice Inspection, page 66-6
- Feature History for the TLS Proxy for Encrypted Voice Inspection, page 66-17

Information about the TLS Proxy for Encrypted Voice Inspection

End-to-end encryption often leaves network security appliances “blind” to media and signaling traffic, which can compromise access control and threat prevention security functions. This lack of visibility can result in a lack of interoperability between the firewall functions and the encrypted voice, leaving businesses unable to satisfy both of their key security requirements.

The ASA is able to intercept and decrypt encrypted signaling from Cisco encrypted endpoints to the Cisco Unified Communications Manager (Cisco UCM), and apply the required threat protection and access control. It can also ensure confidentiality by re-encrypting the traffic onto the Cisco UCM servers.

Typically, the ASA TLS Proxy functionality is deployed in campus unified communications network. This solution is ideal for deployments that utilize end to end encryption and firewalls to protect Unified Communications Manager servers.

The security appliance in Figure 66-1 serves as a proxy for both client and server, with Cisco IP Phone and Cisco UCM interaction.

Figure 66-1 TLS Proxy Flow

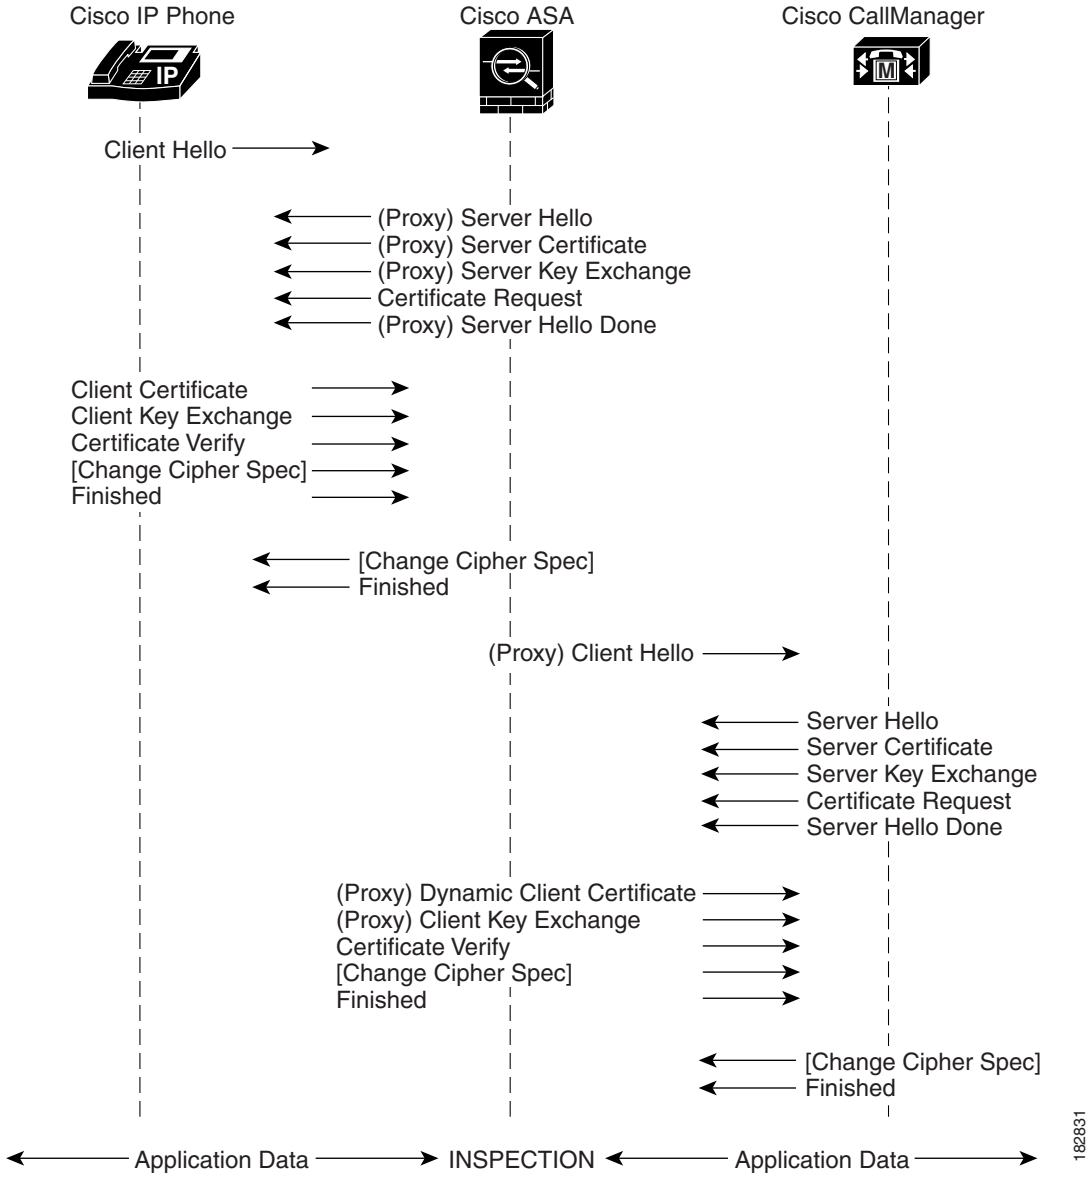

Decryption and Inspection of Unified Communications Encrypted Signaling

With encrypted voice inspection, the security appliance decrypts, inspects and modifies (as needed, for example, performing NAT fixup), and re-encrypts voice signaling traffic while all of the existing VoIP inspection functions for Skinny and SIP protocols are preserved. Once voice signaling is decrypted, the plaintext signaling message is passed to the existing inspection engines.

The security appliance acts as a TLS proxy between the Cisco IP Phone and Cisco UCM. The proxy is transparent for the voice calls between the phone and theCisco UCM. Cisco IP Phones download a Certificate Trust List from the Cisco UCM before registration which contains identities (certificates) of the devices that the phone should trust, such as TFTP servers and Cisco UCM servers. To support server

proxy, the CTL file must contain the certificate that the security appliance creates for the Cisco UCMs. To proxy calls on behalf of the Cisco IP Phone, the security appliance presents a certificate that the Cisco UCM can verify, which is a Local Dynamic Certificate for the phone, issued by the certificate authority on the security appliance.

TLS proxy is supported by the Cisco Unified CallManager Release 5.1 and later. You should be familiar with the security features of the Cisco UCM. For background and detailed description of Cisco UCM security, see the Cisco Unified CallManager document:

http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/5_0/sec_vir/ae/sec504/index.htm

TLS proxy applies to the encryption layer and must be configured with an application layer protocol inspection. You should be familiar with the inspection features on the ASA, especially Skinny and SIP inspection.

Supported Cisco UCM and IP Phones for the TLS Proxy

Cisco Unified Communications Manager

The following releases of the Cisco Unified Communications Manager are supported with the TLS proxy:

- Cisco Unified CallManager Version 4.*x*
- Cisco Unified CallManager Version 5.0
- Cisco Unified CallManager Version 5.1
- Cisco Unified Communications Manager 6.1
- Cisco Unified Communications Manager 7.0
- Cisco Unified Communications Manager 8.0

Cisco Unified IP Phones

The following IP phones in the Cisco Unified IP Phones 7900 Series are supported with the TLS proxy:

- Cisco Unified IP Phone 7985
- Cisco Unified IP Phone 7975
- Cisco Unified IP Phone 7971
- Cisco Unified IP Phone 7970
- Cisco Unified IP Phone 7965
- Cisco Unified IP Phone 7962
- Cisco Unified IP Phone 7961
- Cisco Unified IP Phone 7961G-GE
- Cisco Unified IP Phone 7960
- Cisco Unified IP Phone 7945
- Cisco Unified IP Phone 7942
- Cisco Unified IP Phone 7941
- Cisco Unified IP Phone 7941G-GE
- Cisco Unified IP Phone 7940
- Cisco Unified Wireless IP Phone 7921

- Cisco Unified Wireless IP Phone 7925
- Cisco IP Communicator (CIPC) for softphones

Licensing for the TLS Proxy

The TLS proxy for encrypted voice inspection feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5510	Base License and Security Plus License: 2 sessions. *Optional licenses: 24, 50, or 100 sessions.*
ASA 5520	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5540	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5550	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5580	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASA 5512-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*

Model	License Requirement[1]
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
   - Phone Proxy
   - Presence Federation Proxy
   - Encrypted Voice Inspection

   Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

   Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

   You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

   Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

   Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

   You might also use SRTP encryption sessions for your connections:
   - For K8 licenses, SRTP sessions are limited to 250.
   - For K9 licenses, there is not limit.

   Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.
2. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

Table 66-1 shows the default and maximum TLS session details by platform.

Table 66-1 Default and Maximum TLS Sessions on the Security Appliance

Security Appliance Platform	Default TLS Sessions	Maximum TLS Sessions
ASA 5505	10	80
ASA 5510	100	200
ASA 5520	300	1200
ASA 5540	1000	4500
ASA 5550	2000	4500
ASA 5580	4000	13,000

For more information about licensing, see Chapter 5, "Managing Feature Licenses."

Prerequisites for the TLS Proxy for Encrypted Voice Inspection

Before configuring TLS proxy, the following prerequisites are required:

- You must set clock on the security appliance before configuring TLS proxy. To set the clock manually and display clock, use the **clock set** and **show clock** commands. We recommend that the security appliance use the same NTP server as the Cisco Unified CallManager cluster. TLS handshake may fail due to certificate validation failure if clock is out of sync between the security appliance and the Cisco Unified CallManager server.
- 3DES-AES license is needed to interoperate with the Cisco Unified CallManager. AES is the default cipher used by the Cisco Unified CallManager and Cisco IP Phone.
- Import the following certificates which are stored on the Cisco UCM. These certificates are required by the ASA for the phone proxy.
  - Cisco_Manufacturing_CA
  - CAP-RTP-001
  - CAP-RTP-002
  - CAPF certificate (Optional)

    If LSC provisioning is required or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA.

  See Chapter 65, "Configuring the Cisco Phone Proxy."For example, the CA Manufacturer certificate is required by the phone proxy to validate the IP phone certificate.

Configuring the TLS Proxy for Encrypted Voice Inspection

This section includes the following topics:

- Configure TLS Proxy Pane, page 66-8
- Adding a TLS Proxy Instance, page 66-9
- Add TLS Proxy Instance Wizard – Server Configuration, page 66-9
- Add TLS Proxy Instance Wizard – Client Configuration, page 66-10
- Add TLS Proxy Instance Wizard – Other Steps, page 66-12
- Edit TLS Proxy Instance – Server Configuration, page 66-13
- Edit TLS Proxy Instance – Client Configuration, page 66-14

CTL Provider

Use the CTL Provider option to configure Certificate Trust List provider service.

The CTL Provider pane lets you define and configure Certificate Trust List provider service to enable inspection of encrypted traffic.

Fields

- CTL Provider Name—Lists the CTL Provider name.

- Client Details—Lists the name and IP address of the client.
  - Interface Name—Lists the defined interface name.
  - IP Address—Lists the defined interface IP address.
- Certificate Name—Lists the certificate to be exported.
- Add—Adds a CTL Provider.
- Edit—Edits a CTL Provider.
- Delete—Deletes a CTL Provider.

Add/Edit CTL Provider

The Add/Edit CTL Provider dialog box lets you define the parameters for the CTL Provider.

Fields

- CTL Provider Name—Specifies the CTL Provider name.
- Certificate to be Exported—Specifies the certificate to be exported to the client.
  - Certificate Name—Specifies the name of the certificate to be exported to the client.
  - Manage—Manages identity certificates.
- Client Details—Specifies the clients allowed to connect.
  - Client to be Added—Specifies the client interface and IP address to add to the client list.

    Interface—Specifies client interface.

    IP Address—Specifies the client IP address.

    Add—Adds the new client to the client list.

    Delete—Deletes the selected client from the client list.
- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.
  - Parse the CTL file provided by the CTL Client and install trustpoints—Trustpoints installed by this option have names prefixed with "_internal_CTL_." If disabled, each Call Manager server and CAPF certificate must be manually imported and installed.
  - Port Number—Specifies the port to which the CTL provider listens. The port must be the same as the one listened to by the CallManager servers in the cluster (as configured under Enterprise Parameters on the CallManager administration page). The default is 2444.
  - Authentication—Specifies the username and password that the client authenticates with the provider.

    Username—Client username.

    Password—Client password.

    Confirm Password—Client password.

Configure TLS Proxy Pane

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

You can configure the TLS Proxy from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Configuring a TLS Proxy lets you use the TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and enable the ASA for the Cisco Unified Communications features:

- TLS Proxy for the Cisco Unified Presence Server (CUPS), part of Presence Federation
- TLS Proxy for the Cisco Unified Mobility Advantage (CUMA) server, part of Mobile Advantage
- Phone Proxy

Fields

- TLS Proxy Name—Lists the TLS Proxy name.
- Server Proxy Certificate—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.
- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.
- Client Proxy Certificate—Lists the proxy certificate for the TLS client. The ASA uses the client proxy certificate to authenticate the TLS client during the handshake between the proxy and the TLS client. The certificate can be either self-signed, enrolled with a certificate authority, or issued by the third party.
- Add—Adds a TLS Proxy by launching the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 66-9 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 66-13 and Edit TLS Proxy Instance – Client Configuration, page 66-14.
- 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.

Adding a TLS Proxy Instance

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

This wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1 Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2 To add a new TLS Proxy Instance, click **Add**.

The Add TLS Proxy Instance Wizard opens.

Step 3 In the TLS Proxy Name field, type the TLS Proxy name.

Step 4 Click **Next**.

The Add TLS Proxy Instance Wizard – Server Configuration dialog box opens. In this step of the wizard, configure the server proxy parameters for original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server. See Add TLS Proxy Instance Wizard – Server Configuration, page 66-9.

After configuring the server proxy parameters, the wizard guides you through configuring client proxy parameters (see Add TLS Proxy Instance Wizard – Client Configuration, page 66-10) and provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 66-12).

Add TLS Proxy Instance Wizard – Server Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

The Add TLS Proxy Instance Wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1 Complete the first step of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 66-9.

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 the "Configuring Identity Certificates Authentication" section on page 49-22.

- To select an existing certificate, select one from the drop-down list.

  When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the ASA, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP**_*ctl-instance_filename*.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 3 To install the TLS server certificate in the ASA trust store, so that the ASA can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. See the "Guidelines and Limitations" section on page 49-9. Click **Add** to open the Install Certificate dialog box. See the "Adding or Installing a CA Certificate" section on page 49-11.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 4 To require the ASA to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

Step 5 Click **Next**.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens. In this step of the wizard, configure the client proxy parameters for original TLS Client—the CUMC client for Mobile Advantage, CUP or MS LCS/OCS client for Presence Federation, or the IP phone for the Phone Proxy. See Add TLS Proxy Instance Wizard – Client Configuration, page 66-10.

After configuring the client proxy parameters, the wizard provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 66-12).

Add TLS Proxy Instance Wizard – Client Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

This wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1 Complete the first two steps of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 66-9 and Add TLS Proxy Instance Wizard – Client Configuration, page 66-10.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens.

Step 2 To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

a. Check the Specify the proxy certificate for the TLS Client... check box.

b. Select a certificate from the drop-down list.

Or

To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens. See the "Configuring Identity Certificates Authentication" section on page 49-22.

Note When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 3 To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the ASA acts as the certificate authority and issues certificates to TLS clients.

a. Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens. See the "Configuring Identity Certificates Authentication" section on page 49-22.

Or

Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the ASA. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens. See the "Authenticating Using the Local CA" section on page 49-29.

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 the "Configuring Identity Certificates Authentication" section on page 49-22 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 66-12).

Add TLS Proxy Instance Wizard – Other Steps

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The last dialog box of the Add TLS Proxy Instance Wizard specifies the additional steps required to make TLS Proxy fully functional. In particular, you need to perform the following tasks to complete the TLS Proxy configuration:

- Export the local CA certificate or LDC Issuer and install them on the original TLS server.

  To export the LDC Issuer, go to Configuration > Firewall > Advanced > Certificate Management > Identity Certificates > Export. See the "Exporting an Identity Certificate" section on page 49-25.

- For the TLS Proxy, enable Skinny and SIP inspection between the TLS server and TLS clients. See SIP Inspection, page 59-20 and Skinny (SCCP) Inspection, page 59-31. When you are configuring the TLS Proxy for Presence Federation (which uses CUP), you only enable SIP inspection because the feature supports only the SIP protocol.
- For the TLS Proxy for CUMA, enable MMP inspection.
- When using the internal Certificate Authority of the ASA to sign the LDC Issuer for TLS clients, perform the following:
  - Use the Cisco CTL Client to add the server proxy certificate to the CTL file and install the CTL file on the ASA.

    For information on the Cisco CTL Client, see "Configuring the Cisco CTL Client" in *Cisco Unified CallManager Security Guide*.

    http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/security/5_0_4/secuauth.html

    To install the CTL file on the ASA, go to Configuration > Firewall > Unified Communications > CTL Provider > Add. The Add CTL Provider dialog box opens. For information on using this dialog box to install the CTL file, see Add/Edit CTL Provider, page 66-7.
  - Create a CTL provider instance for connections from the CTL clients. See Add/Edit CTL Provider, page 66-7.

Edit TLS Proxy Instance – Server Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

Use the Edit TLS Proxy – Server Configuration tab to edit the server proxy parameters for the original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server.

Step 1 Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2 To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3 If necessary, click the **Server Configuration** tab.

Step 4 Specify the server proxy certificate by doing one of the following:

- To add a new certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

  When the Phone Proxy is operating in a mixed-mode CUCM cluster, you must import the CUCM certificate by clicking **Add** in the Manage Identify Certificates dialog box. See the "Configuring CA Certificate Authentication" section on page 49-11.

- To select an existing certificate, select one from the drop-down list.

  When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the ASA, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP**_*ctl-instance_filename*.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 5 To install the TLS server certificate in the ASA trust store, so that the ASA can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. See the "Guidelines and Limitations" section on page 49-9. Click **Add** to open the Install Certificate dialog box. See the "Configuring CA Certificate Authentication" section on page 49-11.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 6 To require the ASA to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

Step 7 Click **Apply** to save the changes.

Edit TLS Proxy Instance – Client Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

The fields in the Edit TLS Proxy dialog box are identical to the fields displayed when you add a TLS Proxy instance. Use the Edit TLS Proxy – Client Configuration tab to edit the client proxy parameters for the original TLS Client, such as IP phones, CUMA clients, the Cisco Unified Presence Server (CUPS), or the Microsoft OCS server.

Step 1 Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2 To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3 If necessary, click the **Client Configuration** tab.

Step 4 To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

a. Check the Specify the proxy certificate for the TLS Client... check box.

b. Select a certificate from the drop-down list.

Or

To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens. See the "Configuring Identity Certificates Authentication" section on page 49-22.

Note When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 5 To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the ASA acts as the certificate authority and issues certificates to TLS clients.

a. Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens. See the "Configuring Identity Certificates Authentication" section on page 49-22.

Or

Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the ASA. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens. See the "Authenticating Using the Local CA" section on page 49-29.

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 the "Configuring Identity Certificates Authentication" section on page 49-22 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.

TLS Proxy

This feature is supported only for ASA versions 8.0.x prior to 8.0.4 and for version 8.1.

Note This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

Use the TLS Proxy option to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco CallManager.

The TLS Proxy pane lets you define and configure Transaction Layer Security Proxy to enable inspection of encrypted traffic.

Fields

- TLS Proxy Name—Lists the TLS Proxy name.
- Server—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.
- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.

- Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client or server dynamic certificates.
- Add—Adds a TLS Proxy.
- Edit—Edits a TLS Proxy.
- Delete—Deletes a TLS Proxy.
- Maximum Sessions—Lets you specify the maximum number of TLS Proxy sessions to support.
  - Specify the maximum number of TLS Proxy sessions that the ASA needs to support. By default, ASA supports 300 sessions.—Enables maximum number of sessions option.
  - Maximum number of sessions:—The minimum is 1. The maximum is dependent on the platform. The default is 300.

Add/Edit TLS Proxy

Note This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

The Add/Edit TLS Proxy dialog box lets you define the parameters for the TLS Proxy.

Fields

- TLS Proxy Name—Specifies the TLS Proxy name.
- Server Configuration—Specifies the proxy certificate name.
  - Server—Specifies the trustpoint to be presented during the TLS handshake. The trustpoint could be self-signed or enrolled locally with the certificate service on the proxy.
- Client Configuration—Specifies the local dynamic certificate issuer and key pair.
  - Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.

    Certificate Authority Server—Specifies the certificate authority server.

    Certificate—Specifies a certificate.

    Manage—Configures the local certificate authority. To make configuration changes after it has been configured for the first time, disable the local certificate authority.
  - Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client dynamic certificates.

    Key-Pair Name—Specifies a defined key pair.

    Show—Shows the key pair details, including generation time, usage, modulus size, and key data.

    New—Lets you define a new key pair.
- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.
  - Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

    Add—Adds the selected algorithm to the active list.

    Remove—Removes the selected algorithm from the active list.

– Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and CallManager may be NULL cipher to offload the CallManager.

  Move Up—Moves an algorithm up in the list.

  Move Down—Moves an algorithm down in the list.

Feature History for the TLS Proxy for Encrypted Voice Inspection

Table 66-2 lists the release history for this feature.

Table 66-2 Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
TLS Proxy	8.0(2)	The TLS proxy feature was introduced.

Configuring Cisco Mobility Advantage

This chapter describes how to configure the adaptive security appliance for Cisco Unified Communications Mobility Advantage Proxy features.

This chapter includes the following sections:

- Information about the Cisco Mobility Advantage Proxy Feature, page 67-1
- Licensing for the Cisco Mobility Advantage Proxy Feature, page 67-6
- Configuring Cisco Mobility Advantage, page 67-6
- Feature History for Cisco Mobility Advantage, page 67-7

Information about the Cisco Mobility Advantage Proxy Feature

This section contains the following topics:

- Cisco Mobility Advantage Proxy Functionality, page 67-1
- Mobility Advantage Proxy Deployment Scenarios, page 67-2
- Trust Relationships for Cisco UMA Deployments, page 67-4

Cisco Mobility Advantage Proxy Functionality

To support Cisco UMA for the Cisco Mobility Advantage solution, the mobility advantage proxy (implemented as a TLS proxy) includes the following functionality:

- The ability to allow no client authentication during the handshake with clients.
- Allowing an imported PKCS-12 certificate to server as a proxy certificate.

The ASA includes an inspection engine to validate the Cisco UMA Mobile Multiplexing Protocol (MMP).

MMP is a data transport protocol for transmitting data entities between Cisco UMA clients and servers. As shown in Figure 67-1, MMP must be run on top of a connection-oriented protocol (the underlying transport) and is intended to be run on top of a secure transport protocol such as TLS. The Orative Markup Language (OML) protocol is intended to be run on top of MMP for the purposes of data synchronization, as well as the HTTP protocol for uploading and downloading large files.

Figure 67-1 MMP Stack

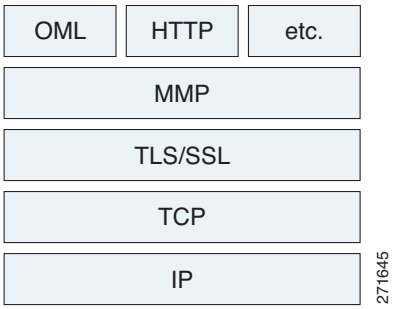

The TCP/TLS default port is 5443. There are no embedded NAT or secondary connections.

Cisco UMA client and server communications can be proxied via TLS, which decrypts the data, passes it to the inspect MMP module, and re-encrypt the data before forwarding it to the endpoint. The inspect MMP module verifies the integrity of the MMP headers and passes the OML/HTTP to an appropriate handler. The ASA takes the following actions on the MMP headers and data:

- Verifies that client MMP headers are well-formed. Upon detection of a malformed header, the TCP session is terminated.
- Verifies that client to server MMP header lengths are not exceeded. If an MMP header length is exceeded (4096), then the TCP session is terminated.
- Verifies that client to server MMP content lengths are not exceeded. If an entity content length is exceeded (4096), the TCP session is terminated.

Note 4096 is the value currently used in MMP implementations.

Because MMP headers and entities can be split across packets, the ASA buffers data to ensure consistent inspection. The SAPI (stream API) handles data buffering for pending inspection opportunities. MMP header text is treated as case insensitive and a space is present between header text and values. Reclaiming of MMP state is performed by monitoring the state of the TCP connection.

Mobility Advantage Proxy Deployment Scenarios

Figure 67-2 and Figure 67-3 show the two deployment scenarios for the TLS proxy used by the Cisco Mobility Advantage solution. In scenario 1 (the recommended deployment architecture), the ASA functions as both the firewall and TLS proxy. In scenario 2, the ASA functions as the TLS proxy only and works with an existing firewall. In both scenarios, the clients connect from the Internet.

In the scenario 1 deployment, the ASA is between a Cisco UMA client and a Cisco UMA server. The Cisco UMA client is an executable that is downloaded to each smartphone. The Cisco UMA client applications establishes a data connection, which is a TLS connection, to the corporate Cisco UMA server. The ASA intercepts the connections and inspects the data that the client sends to the Cisco UMA server.

Figure 67-2 The TLS proxy for the Cisco Mobility Advantage solution does not support client authentication because the Cisco UMA client cannot present a certificate. ***Security Appliance as Firewall with Mobility Advantage Proxy and MMP Inspection***

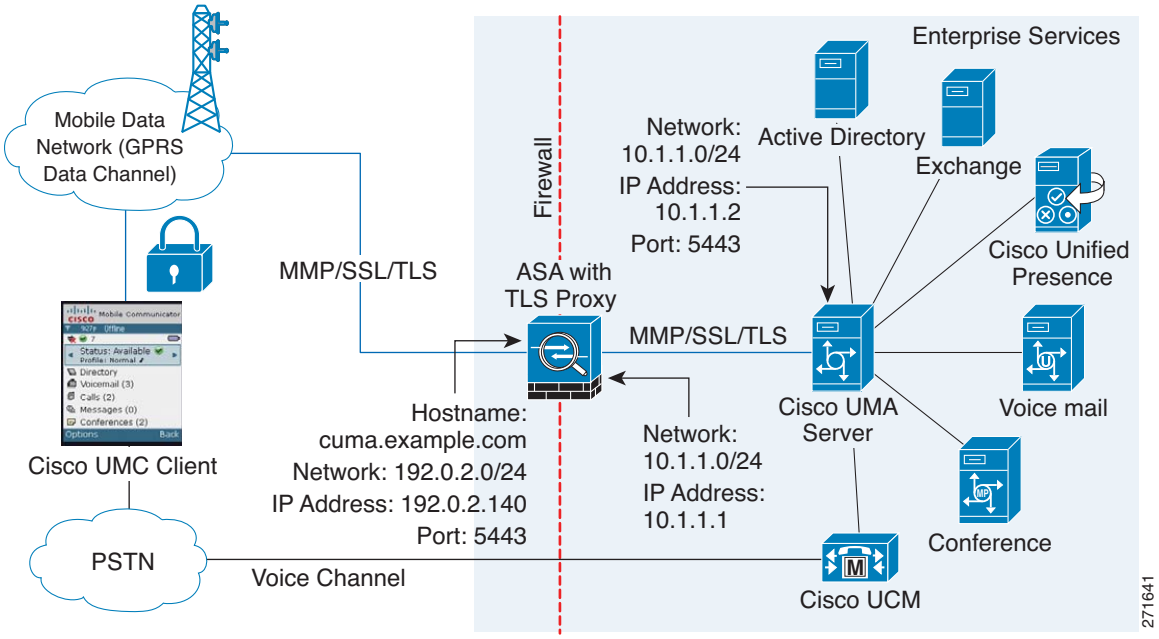

In Figure 67-2, the ASA performs static NAT by translating the Cisco UMA server 10.1.1.2 IP address to 192.0.2.140.

Figure 67-3 shows deployment scenario 2, where the ASA functions as the TLS proxy only and does not function as the corporate firewall. In this scenario, the ASA and the corporate firewall are performing NAT. The corporate firewall will not be able to predict which client from the Internet needs to connect to the corporate Cisco UMA server. Therefore, to support this deployment, you can take the following actions:

- Set up a NAT rule for inbound traffic that translates the destination IP address 192.0.2.41 to 172.16.27.41.
- Set up an interface PAT rule for inbound traffic translating the source IP address of every packet so that the corporate firewall does not need to open up a wildcard pinhole. The Cisco UMA server receives packets with the source IP address 192.0.12.183.

  See Chapter 40, "Configuring Network Object NAT (ASA 8.3 and Later)" and Chapter 41, "Configuring Twice NAT (ASA 8.3 and Later)" for information.

Note This interface PAT rule converges the Cisco UMA client IP addresses on the outside interface of the ASA into a single IP address on the inside interface by using different source ports. Performing this action is often referred as "outside PAT". "Outside PAT" is not recommended when TLS proxy for Cisco Mobility Advantage is enabled on the same interface of the ASA with phone proxy, Cisco Unified Presence, or any other features involving application inspection. "Outside PAT" is not supported completely by application inspection when embedded address translation is needed.

Figure 67-3 Cisco UMC/Cisco UMA Architecture – Scenario 2: Security Appliance as Mobility Advantage Proxy Only

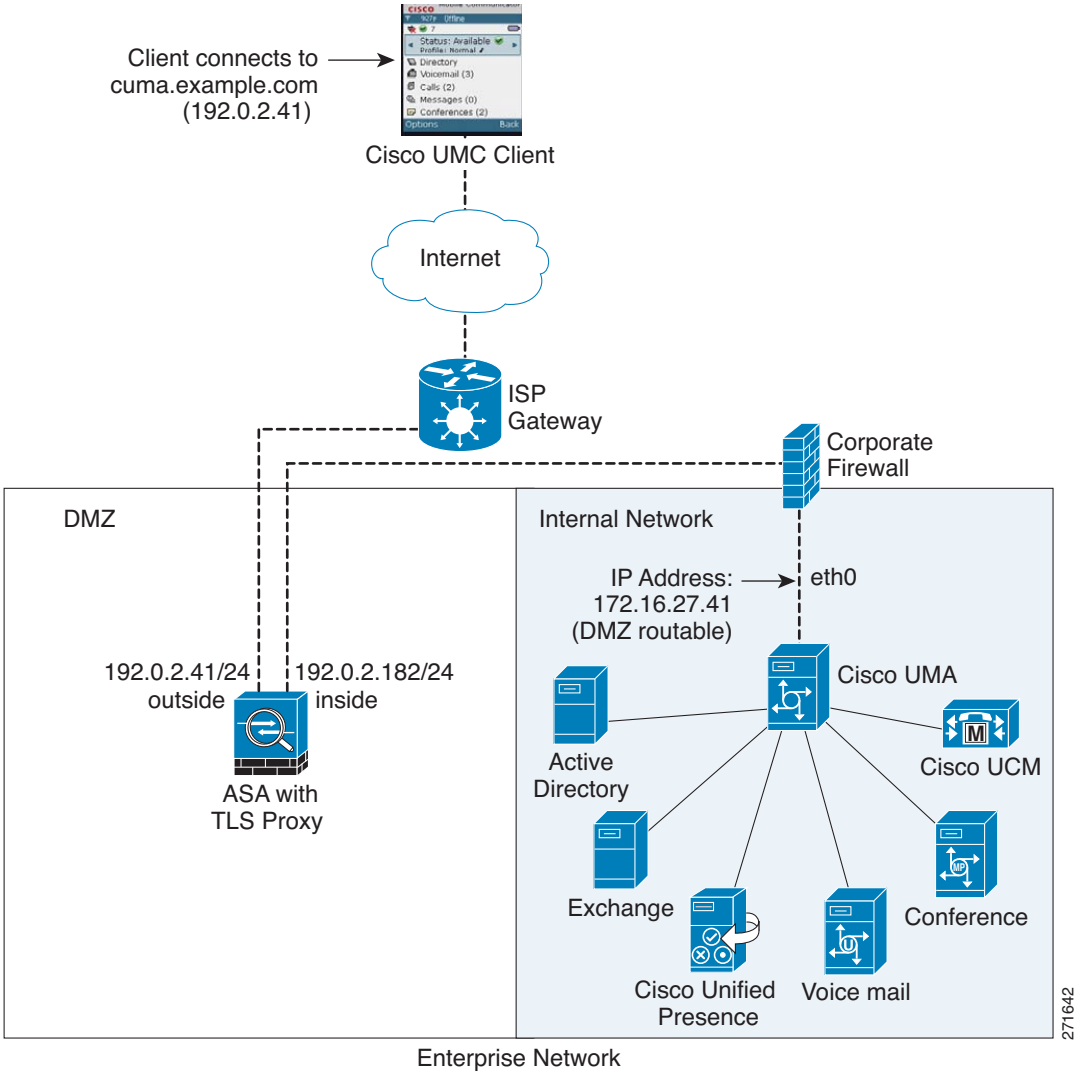

Mobility Advantage Proxy Using NAT/PAT

In both scenarios (Figure 67-2 and Figure 67-3), NAT can be used to hide the private address of the Cisco UMA servers.

In scenario 2 (Figure 67-3), PAT can be used to converge all client traffic into one source IP, so that the firewall does not have to open up a wildcard pinhole for inbound traffic.

Trust Relationships for Cisco UMA Deployments

To establish a trust relationship between the Cisco UMC client and the ASA, the ASA uses the Cisco UMA server certificate and keypair or the ASA obtains a certificate with the Cisco UMA server FQDN (certificate impersonation). Between the ASA and the Cisco UMA server, the ASA and Cisco UMA server use self-signed certificates or certificates issued by a local certificate authority.

Figure 67-4 shows how you can import the Cisco UMA server certificate onto the ASA. When the Cisco UMA server has already enrolled with a third-party CA, you can import the certificate with the private key onto the ASA. Then, the ASA has the full credentials of the Cisco UMA server. When a Cisco UMA client connects to the Cisco UMA server, the ASA intercepts the handshake and uses the Cisco UMA server certificate to perform the handshake with the client. The ASA also performs a handshake with the server.

Figure 67-4 How the Security Appliance Represents Cisco UMA – Private Key Sharing

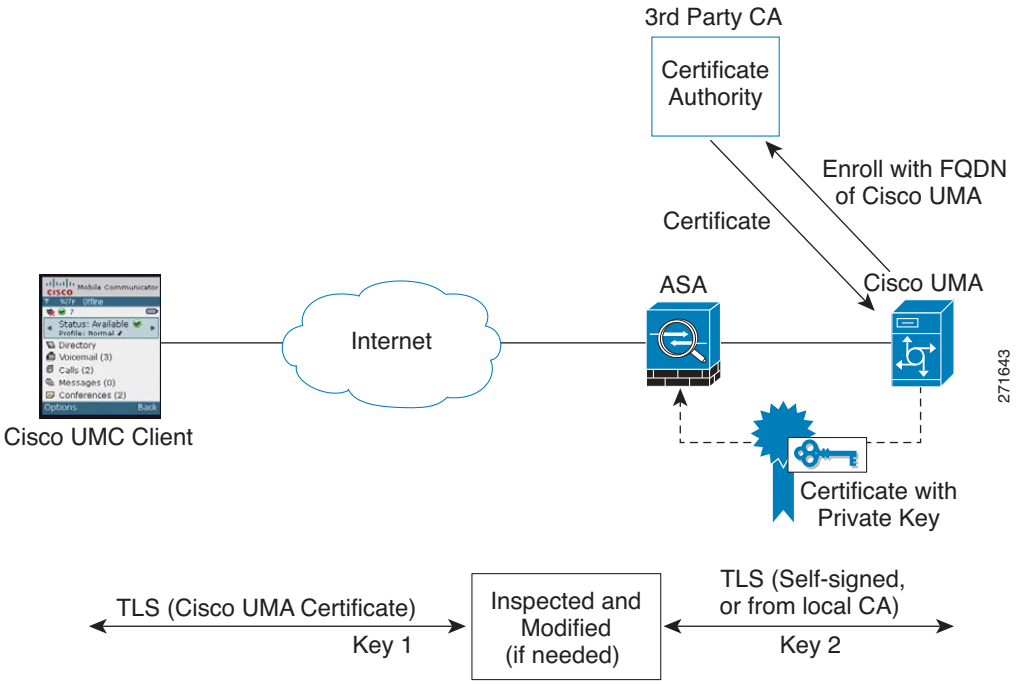

Figure 67-5 shows another way to establish the trust relationship. Figure 67-5 shows a green field deployment, because each component of the deployment has been newly installed. The ASA enrolls with the third-party CA by using the Cisco UMA server FQDN as if the ASA is the Cisco UMA server. When the Cisco UMA client connects to the ASA, the ASA presents the certificate that has the Cisco UMA server FQDN. The Cisco UMA client believes it is communicating to with the Cisco UMA server.

Figure 67-5 How the Security Appliance Represents Cisco UMA – Certificate Impersonation

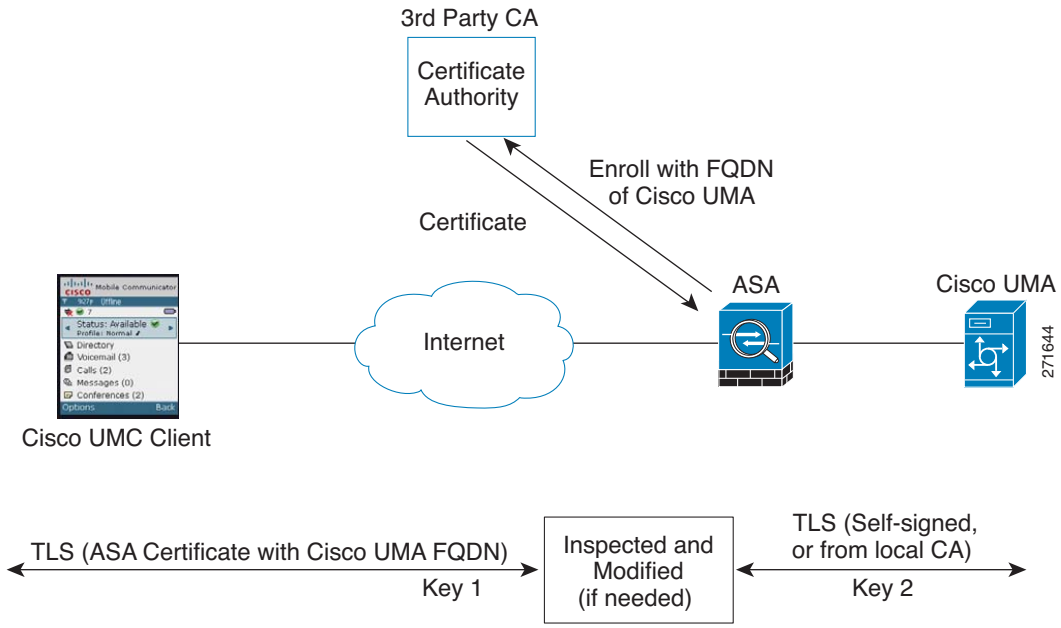

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The ASA's identity certificate is exported, and then uploaded on the Cisco UMA server truststore. The Cisco UMA server certificate is downloaded, and then uploaded on the ASA truststore by creating a trustpoint and using the **crypto ca authenticate** command.

Licensing for the Cisco Mobility Advantage Proxy Feature

The Cisco Unified Communications proxy features (Cisco Phone Proxy, TLS proxy for encrypted voice inspection, and the Cisco Presence Federation Proxy) supported by the ASA require a Unified Communications Proxy license. However, in Version 8.2(2) and later, the Mobility Advantage proxy no longer requires a Unified Communications Proxy license.

The following table shows the licensing requirements for the Mobility Advantage proxy:

Model	License Requirement
All models	Base License.

For more information about licensing, see Chapter 5, “Managing Feature Licenses.”

Configuring Cisco Mobility Advantage

This section includes the following topic:

- Task Flow for Configuring Cisco Mobility Advantage, page 67-7

Task Flow for Configuring Cisco Mobility Advantage

To configure for the ASA to perform TLS proxy and MMP inspection as shown in Figure 67-2 and Figure 67-3, perform the following tasks.

It is assumed that self-signed certificates are used between the ASA and the Cisco UMA server.

To configure the Cisco Mobility Advantage Proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Mobility Advantage Proxy option under the Remote Access section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Unified Presence Proxy instance, importing and installing the required certificates, and finally enables the MMP inspection for the Mobility Advantage traffic automatically.

The wizard guides you through four steps to create the Mobility Advantage Proxy:

Step 1 Select the Mobility Advantage Proxy option.

Step 2 Specify setting to define the proxy topology, such the IP address of the Mobility Advantage server.

Step 3 Configure the server-side certificate management, namely the certificates that are exchanged between the local Mobility Advantage server and the ASA.

Step 4 Configure the client-side certificate management, namely the certificates that are exchanged between the Unified Mobile Communicator and the ASA

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy. See Chapter 64, "Using the Cisco Unified Communication Wizard" for more information.

Feature History for Cisco Mobility Advantage

Table 67-1 lists the release history for this feature.

Table 67-1 Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Mobility Advantage Proxy	8.0(4)	The Cisco Mobility Advantage Proxy feature was introduced.
Cisco Mobility Advantage Proxy	8.3(1)	The Unified Communications Wizard was added to ASDM. By using the wizard, you can configure the Cisco Mobility Advantage Proxy.

Configuring Cisco Unified Presence

This chapter describes how to configure the adaptive security appliance for Cisco Unified Presence.

This chapter includes the following sections:

- Information About Cisco Unified Presence, page 68-1
- Licensing for Cisco Unified Presence, page 68-7
- Configuring Cisco Unified Presence Proxy for SIP Federation, page 68-8
- Feature History for Cisco Unified Presence, page 68-9

Information About Cisco Unified Presence

This section includes the following topics:

- Architecture for Cisco Unified Presence for SIP Federation Deployments, page 68-1
- Trust Relationship in the Presence Federation, page 68-4
- Security Certificate Exchange Between Cisco UP and the Security Appliance, page 68-5
- XMPP Federation Deployments, page 68-5
- Configuration Requirements for XMPP Federation, page 68-6

Architecture for Cisco Unified Presence for SIP Federation Deployments

Figure 68-1 depicts a Cisco Unified Presence/LCS Federation scenario with the ASA as the presence federation proxy (implemented as a TLS proxy). The two entities with a TLS connection are the "Routing Proxy" (a dedicated Cisco UP) in Enterprise X and the Microsoft Access Proxy in Enterprise Y. However, the deployment is not limited to this scenario. Any Cisco UP or Cisco UP cluster could be deployed on the left side of the ASA; the remote entity could be any server (an LCS, an OCS, or another Cisco UP).

The following architecture is generic for two servers using SIP (or other ASA inspected protocols) with a TLS connection.

Entity X: Cisco UP/Routing Proxy in Enterprise X

Entity Y: Microsoft Access Proxy/Edge server for LCS/OCS in Enterprise Y

Figure 68-1 Typical Cisco Unified Presence/LCS Federation Scenario

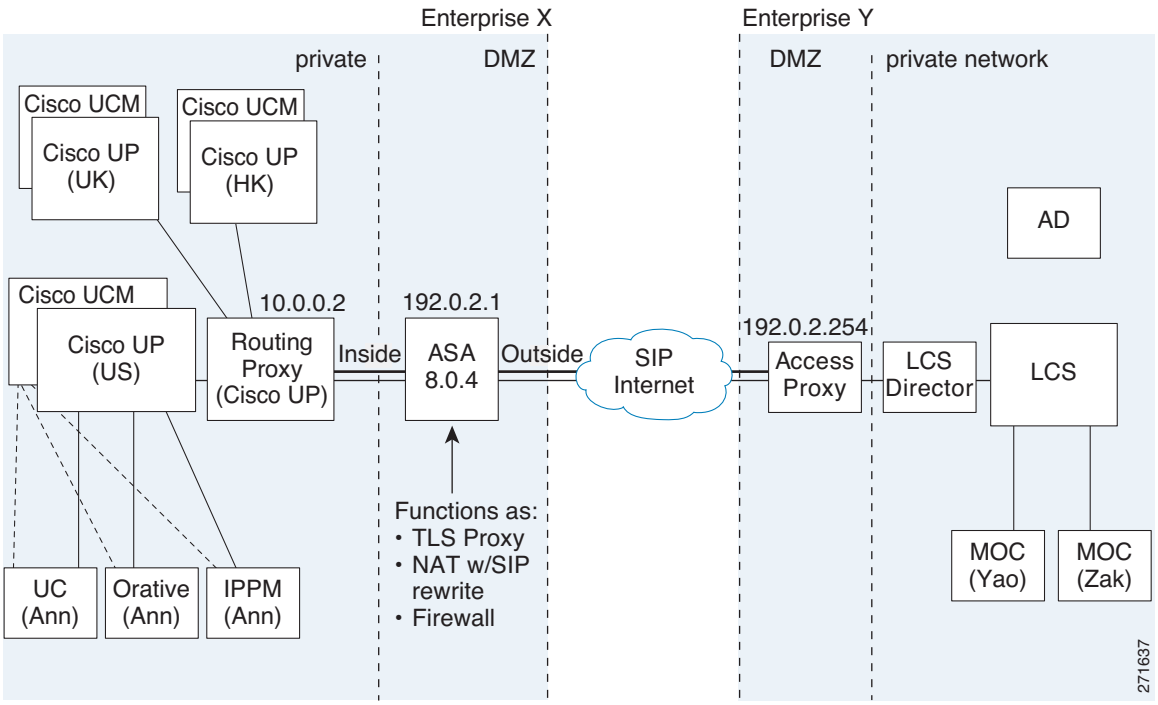

In the above architecture, the ASA functions as a firewall, NAT, and TLS proxy, which is the recommended architecture. However, the ASA can also function as NAT and the TLS proxy alone, working with an existing firewall.

Either server can initiate the TLS handshake (unlike IP Telephony or Cisco Unified Mobility, where only the clients initiate the TLS handshake). There are by-directional TLS proxy rules and configuration. Each enterprise can have an ASA as the TLS proxy.

In Figure 68-1, NAT or PAT can be used to hide the private address of Entity X. In this situation, static NAT or PAT must be configured for foreign server (Entity Y) initiated connections or the TLS handshake (inbound). Typically, the public port should be 5061. The following static PAT command is required for the Cisco UP that accepts inbound connections:

```
hostname(config)# object network obj-10.0.0.2-01
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5061
5061
```

The following static PAT must be configured for each Cisco UP that could initiate a connection (by sending SIP SUBSCRIBE) to the foreign server.

For Cisco UP with the address 10.0.0.2, enter the following command:

```
hostname(config)# object network obj-10.0.0.2-02
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5062
5062
hostname(config)# object network obj-10.0.0.2-03
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service udp 5070
5070
hostname(config)# object network obj-10.0.0.2-04
hostname(config-network-object)# host 10.0.0.2
```

```
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5060
5060
```

For another Cisco UP with the address 10.0.0.3, you must use a different set of PAT ports, such as 45062 or 45070:

```
hostname(config)# object network obj-10.0.0.3-01
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5061
45061
hostname(config)# object network obj-10.0.0.3-02
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5062
45062
hostname(config)# object network obj-10.0.0.3-03
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service udp 5070
5070
hostname(config)# object network obj-10.0.0.2-03
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5070
45070
hostname(config)# object network obj-10.0.0.3-04
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5060
45060
```

Dynamic NAT or PAT can be used for the rest of the outbound connections or the TLS handshake. The ASA SIP inspection engine takes care of the necessary translation (fixup).

```
hostname(config)# object network obj-0.0.0.0-01
hostname(config-network-object)# subnet 0.0.0.0 0.0.0.0
hostname(config-network-object)# nat (inside,outside) dynamic 192.0.2.1
```

Figure 68-2 illustrates an abstracted scenario with Entity X connected to Entity Y through the presence federation proxy on the ASA. The proxy is in the same administrative domain as Entity X. Entity Y could have another ASA as the proxy but this is omitted for simplicity.

Figure 68-2 Abstracted Presence Federation Proxy Scenario between Two Server Entities

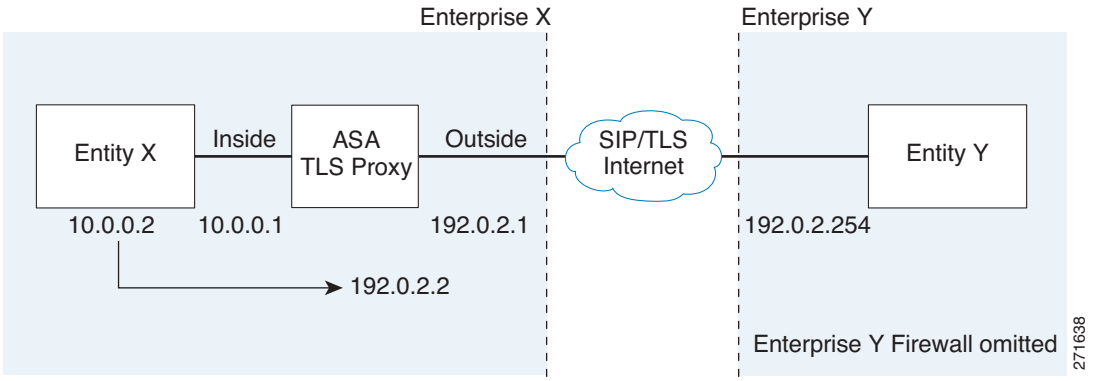

For the Entity X domain name to be resolved correctly when the ASA holds its credential, the ASA could be configured to perform NAT for Entity X, and the domain name is resolved as the Entity X public address for which the ASA provides proxy service.

For further information about configuring Cisco Unified Presence Federation for SIP Federation, see the Integration Guide for Configuring Cisco Unified Presence for Interdomain Federation.:

http://www.cisco.com/en/US/products/ps6837/products_installation_and_configuration_guides_list.html

Trust Relationship in the Presence Federation

Within an enterprise, setting up a trust relationship is achievable by using self-signed certificates or you can set it up on an internal CA.

Establishing a trust relationship cross enterprises or across administrative domains is key for federation. Cross enterprises you must use a trusted third-party CA (such as, VeriSign). The ASA obtains a certificate with the FQDN of the Cisco UP (certificate impersonation).

For the TLS handshake, the two entities could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. Both entities enroll with the CAs. The ASA as the TLS proxy must be trusted by both entities. The ASA is always associated with one of the enterprises. Within that enterprise (Enterprise X in Figure 68-1), the entity and the ASA could authenticate each other via a local CA, or by using self-signed certificates.

To establish a trusted relationship between the ASA and the remote entity (Entity Y), the ASA can enroll with the CA on behalf of Entity X (Cisco UP). In the enrollment request, the Entity X identity (domain name) is used.

Figure 68-3 shows the way to establish the trust relationship. The ASA enrolls with the third party CA by using the Cisco UP FQDN as if the ASA is the Cisco UP.

Figure 68-3 How the Security Appliance Represents Cisco Unified Presence – Certificate Impersonate

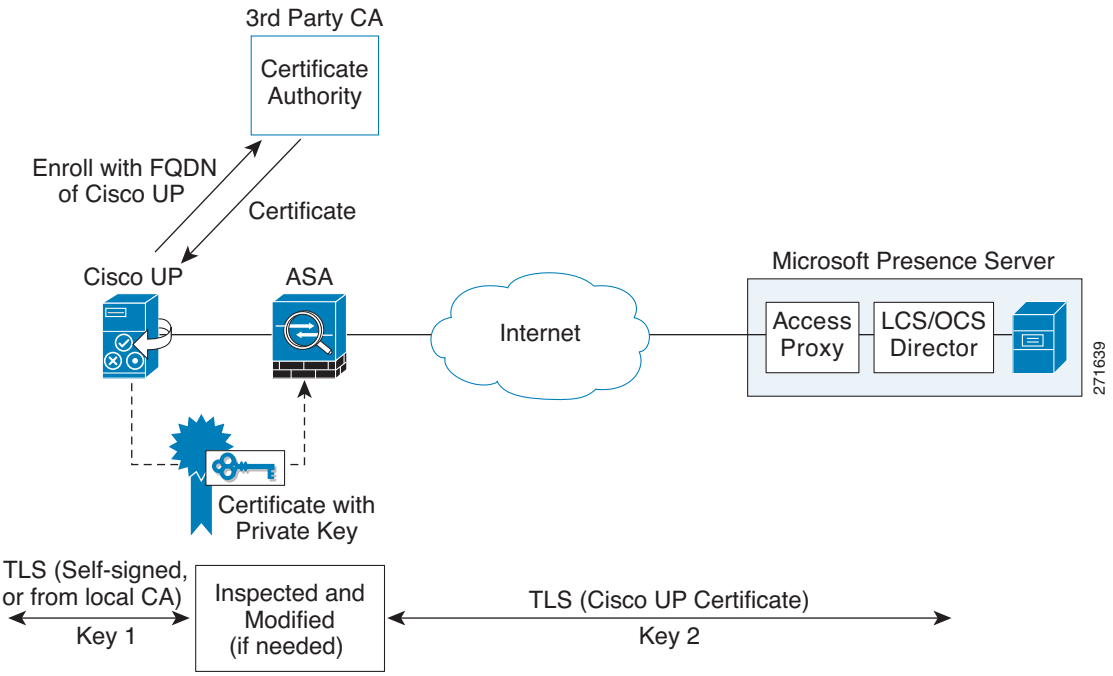

Security Certificate Exchange Between Cisco UP and the Security Appliance

You need to generate the keypair for the certificate (such as `cup_proxy_key`) used by the ASA, and configure a trustpoint to identify the self-signed certificate sent by the ASA to Cisco UP (such as `cup_proxy`) in the TLS handshake.

For the ASA to trust the Cisco UP certificate, you need to create a trustpoint to identify the certificate from the Cisco UP (such as `cert_from_cup`), and specify the enrollment type as terminal to indicate that you will paste the certificate received from the Cisco UP into the terminal.

XMPP Federation Deployments

Figure 68-4 provides an example of an XMPP federated network between Cisco Unified Presence enterprise deployment and an IBM Sametime enterprise deployment. TLS is optional for XMPP federation. ASA acts only as a firewall for XMPP federation; it does not provide TLS proxy functionality or PAT for XMPP federation.

Figure 68-4 Basic XMPP Federated Network between Cisco Unified Presence and IBM Sametime

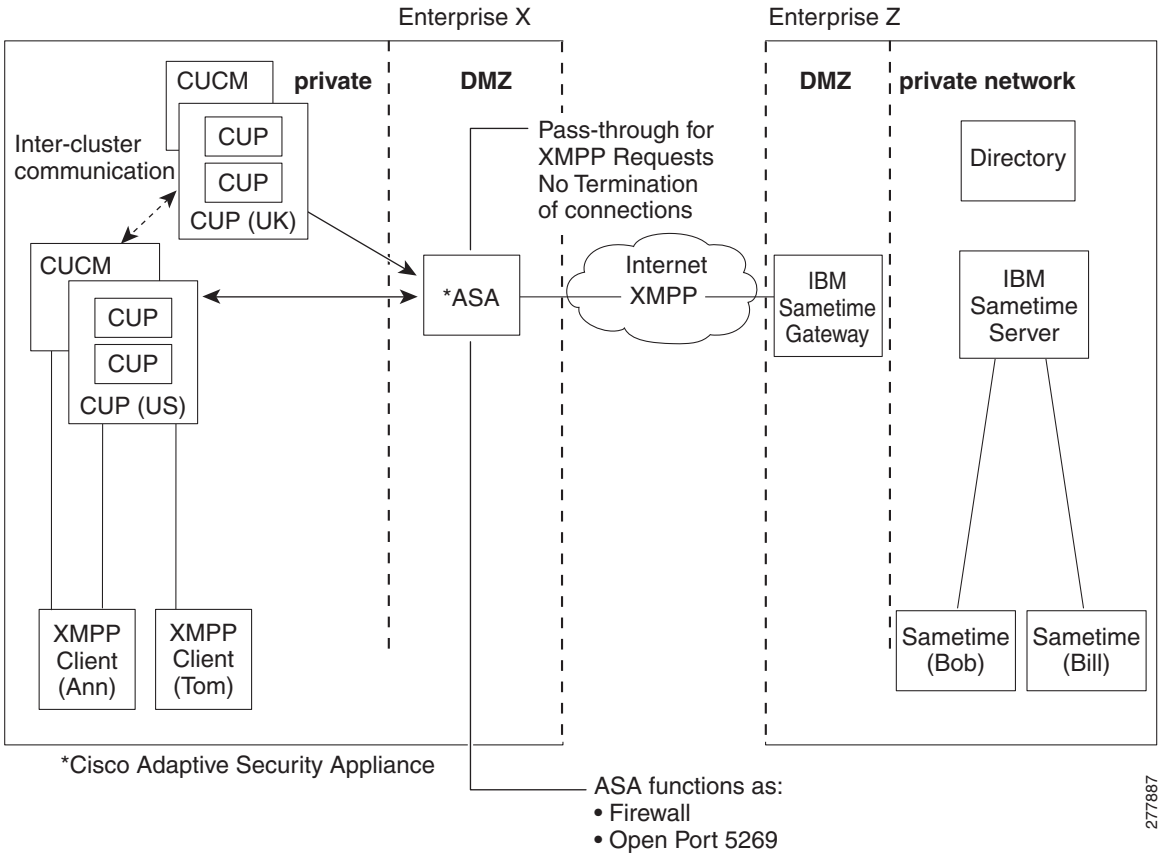

There are two DNS servers within the internal Cisco Unified Presence enterprise deployment. One DNS server hosts the Cisco Unified Presence private address. The other DNS server hosts the Cisco Unified Presence public address and a DNS SRV records for SIP federation (_sipfederationtle), and XMPP federation (_xmpp-server) with Cisco Unified Presence. The DNS server that hosts the Cisco Unified Presence public address is located in the local DMZ.

For further information about configuring Cisco Unified Presence Federation for XMPP Federation, see the *Integration Guide for Configuring Cisco Unified Presence Release 8.0 for Interdomain Federation*:

http://www.cisco.com/en/US/products/ps6837/products_installation_and_configuration_guides_list.html

Configuration Requirements for XMPP Federation

For XMPP Federation, ASA acts as a firewall only. You must open port 5269 for both incoming and outgoing XMPP federated traffic on ASA.

These are sample access lists to open port 5269 on ASA.

Allow traffic from any address to any address on port 5269:

```
access-list ALLOW-ALL extended permit tcp any any eq 5269
```

Allow traffic from any address to any single node on port 5269:

```
access-list ALLOW-ALL extended permit tcp any host <private cup IP address> eq 5269
```

If you do not configure the access list above, and you publish additional XMPP federation nodes in DNS, you must configure access to each of these nodes, for example:

```
object network obj_host_<private cup ip address>
#host <private cup ip address>
object network obj_host_<private cup2 ip address>
#host <private cup2 ip address>
object network obj_host_<public cup ip address>
#host <public cup ip address>
....
```

Configure the following NAT commands:

```
nat (inside,outside) source static obj_host_<private cup1 IP> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup1 IP> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269
```

If you publish a single public IP address in DNS, and use arbitrary ports, configure the following:

(This example is for two additional XMPP federation nodes)

```
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_25269
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_25269

nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_35269
nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_35269
```

If you publish multiple public IP addresses in DNS all using port 5269, configure the following:

(This example is for two additional XMPP federation nodes)

```
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup2 IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup2 IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269

nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup3 IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269
```

Licensing for Cisco Unified Presence

The Cisco Unified Presence feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5510	Base License and Security Plus License: 2 sessions. *Optional licenses: 24, 50, or 100 sessions.*
ASA 5520	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5540	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5550	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5580	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASA 5512-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*

Model	License Requirement[1]
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*[2]

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
   - Phone Proxy
   - Presence Federation Proxy
   - Encrypted Voice Inspection

   Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

   Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

   You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

   Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

   Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

   You might also use SRTP encryption sessions for your connections:
   - For K8 licenses, SRTP sessions are limited to 250.
   - For K9 licenses, there is not limit.

   Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.
2. With the 10,000-session UC license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

For more information about licensing, see Chapter 5, "Managing Feature Licenses."

Configuring Cisco Unified Presence Proxy for SIP Federation

This section contains the following topic:

- Task Flow for Configuring Cisco Unified Presence Federation Proxy for SIP Federation, page 68-9

Task Flow for Configuring Cisco Unified Presence Federation Proxy for SIP Federation

To configure a Cisco Unified Presence/LCS Federation scenario with the ASA as the TLS proxy where there is a single Cisco UP that is in the local domain and self-signed certificates are used between the Cisco UP and the ASA (like the scenario shown in Figure 68-1), perform the following tasks.

To configure the Cisco Unified Presence proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Unified Presence Proxy option under the Business-to-Business section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Unified Presence Proxy instance, importing and installing the required certificates, and finally enables the SIP and SCCP inspection for the Presence Federation traffic automatically.

The wizard guides you through four steps to create the Presence Federation Proxy:

Step 1 Select the Presence Federation Proxy option.

Step 2 Specify setting to define the proxy topology, such the IP address of the Presence Federation server.

Step 3 Configure the local-side certificate management, namely the certificates that are exchanged between the local Unified Presence Federation server and the ASA.

Step 4 Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA

The wizard completes by displaying a summary of the configuration created for Presence Federation. See the Unified Communications Wizard section in this documentation for more information.

Feature History for Cisco Unified Presence

Table 68-1 lists the release history for this feature.

Table 68-1 Feature History for Cisco Unified Presence

Feature Name	Releases	Feature Information
Cisco Presence Federation Proxy	8.0(4)	The Cisco Unified Presence proxy feature was introduced.
Cisco Presence Federation Proxy	8.3(1)	The Unified Communications Wizard was added to ASDM. By using the wizard, you can configure the Cisco Presence Federation Proxy. Support for XMPP Federation was introduced.

Configuring Cisco Intercompany Media Engine Proxy

This chapter describes how to configure the adaptive security appliance for Cisco Intercompany Media Engine Proxy.

This chapter includes the following sections:

- Information About Cisco Intercompany Media Engine Proxy, page 69-1
- Licensing for Cisco Intercompany Media Engine, page 69-8
- Guidelines and Limitations, page 69-9
- Configuring Cisco Intercompany Media Engine Proxy, page 69-11
- Feature History for Cisco Intercompany Media Engine Proxy, page 69-37

Information About Cisco Intercompany Media Engine Proxy

This section includes the following topics:

- Features of Cisco Intercompany Media Engine Proxy, page 69-1
- How the UC-IME Works with the PSTN and the Internet, page 69-2
- Tickets and Passwords, page 69-3
- Call Fallback to the PSTN, page 69-5
- Architecture and Deployment Scenarios for Cisco Intercompany Media Engine, page 69-5

Features of Cisco Intercompany Media Engine Proxy

Cisco Intercompany Media Engine enables companies to interconnect on-demand, over the Internet with advanced features made available by VoIP technologies. Cisco Intercompany Media Engine allows for business-to-business federation between Cisco Unified Communications Manager clusters in different enterprises by utilizing peer-to-peer, security, and SIP protocols to create dynamic SIP trunks between businesses. A collection of enterprises work together to end up looking like one large business with inter-cluster trunks between them.

The adaptive security appliance applies its existing TLS proxy, SIP Application Layer Gateway (ALG), and SIP verification features to the functioning of Cisco Intercompany Media Engine.

Cisco Intercompany Media Engine has the following key features:

- Works with existing phone numbers: Cisco Intercompany Media Engine works with the phone numbers an enterprise currently has and does not require an enterprise to learn new numbers or change providers to use Cisco Intercompany Media Engine.
- Works with existing IP phones: Cisco Intercompany Media Engine works with the existing IP phones within an enterprise. However, the feature set in business-to-business calls is limited to the capabilities of the IP phones.
- Does not require purchasing new services: Cisco Intercompany Media Engine does not require any new services from any service providers. Customers continue to use the PSTN connectivity they have and the Internet connectivity they have today. Cisco Intercompany Media Engine gradually moves calls off the PSTN and onto the Internet.
- Provides a full Cisco Unified Communications experience: Because Cisco Intercompany Media Engine creates inter-cluster SIP trunks between enterprises, any Unified Communication features that work over the SIP trunk and only require a SIP trunk work with the Cisco Intercompany Media Engine, thus providing a Unified Communication experience across enterprises.
- Works on the Internet: Cisco Intercompany Media Engine was designed to work on the Internet. It can also work on managed extranets.
- Provides worldwide reach: Cisco Intercompany Media Engine can connect to any enterprise anywhere in the world, as long as the enterprise is running Cisco Intercompany Media Engine technology. There are no regional limitations. This is because Cisco Intercompany Media Engine utilizes two networks that both have worldwide reach—the Internet and the PSTN.
- Allows for unlimited scale: Cisco Intercompany Media Engine can work with any number of enterprises.
- Is self-learning: The system is primarily self-learning. Customers do not have to enter information about other businesses: no phone prefixes, no IP address, no ports, no domain names, nor certificates. Customers need to configure information about their own networks, and provide policy information if they want to limit the scope of Cisco Intercompany Media Engine.
- Is secure: Cisco Intercompany Media Engine is secure, utilizing a large number of different technologies to accomplish this security.
- Includes anti-spam: Cisco Intercompany Media Engine prevents people from setting up software on the Internet that spams enterprises with phone calls. It provides an extremely high barrier to entry.
- Provides for QoS management: Cisco Intercompany Media Engine provides features that help customers manage the QoS on the Internet, such as the ability to monitor QoS of the RTP traffic in real-time and fallback to PSTN automatically if problems arise.

How the UC-IME Works with the PSTN and the Internet

The Cisco Intercompany Media Engine utilizes two networks that both have worldwide reach—the Internet and the PSTN. Customers continue to use the PSTN connectivity they have. The Cisco Intercompany Media Engine gradually moves calls off the PSTN and onto the Internet. However, if QoS problems arise, the Cisco Intercompany Media Engine Proxy monitors QoS of the RTP traffic in real-time and fallbacks to PSTN automatically.

The Cisco Intercompany Media Engine uses information from PSTN calls to validate that the terminating side owns the number that the originated side had called. After the PSTN call terminates, the enterprises involved in the call send information about the call to their Cisco IME server. The Cisco IME server on the originating side validates the call. Figure 69-1 shows the initial call flow through the PSTN.

On successful verification, the terminating side creates a ticket that grants permission to the call originator to make a Cisco IME call to a specific number. See Tickets and Passwords, page 69-3 for information.

Figure 69-1 Interaction of the UC-IME Proxy with the PSTN

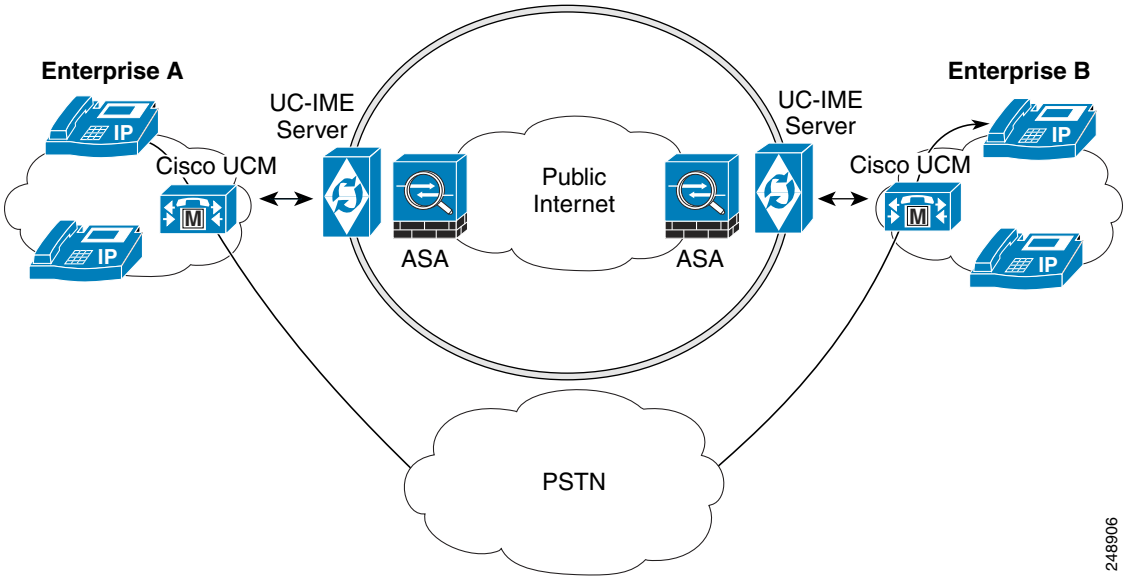

Tickets and Passwords

Cisco Intercompany Media Engine utilizes tickets and passwords to provide enterprise verification. Verification through the creation of tickets ensures an enterprise is not subject to denial-of-service (DOS) attacks from the Internet or endless VoIP spam calls. Ticket verification prevents spam and DOS attacks because it introduces a cost to the VoIP caller; namely, the cost of a PSTN call. A malicious user cannot set up just an open source asterisk PBX on the Internet and begin launching SIP calls into an enterprise running Cisco Intercompany Media Engine. Having the Cisco Intercompany Media Engine Proxy verify tickets allows incoming calls from a particular enterprise to a particular number only when that particular enterprise has previously called that phone number on the PSTN.

To send a spam VoIP call to every phone within an enterprise, an organization would have to purchase the Cisco Intercompany Media Engine and Cisco Unified Communications Manager and have called each phone number within the enterprise over the PSTN and completed each call successfully. Only then can it launch a VoIP call to each number.

The Cisco Intercompany Media Engine server creates tickets and the ASA validates them. The ASA and Cisco Intercompany Media Engine server share a password that is configured so that the ASA detects the ticket was created by a trusted Cisco Intercompany Media Engine server. The ticket contains information that indicates that the enterprise is authorized to call specific phone numbers at the target enterprise. See Figure 69-2 for the ticket verification process and how it operates between the originating and terminating-call enterprises.

Note Because the initial calls are over the PSTN, they are subject to any national regulations regarding telemarketing calling. For example, within the United States, they would be subject to the national do-not-call registry.

Figure 69-2 Ticket Verification Process with Cisco Intercompany Media Engine

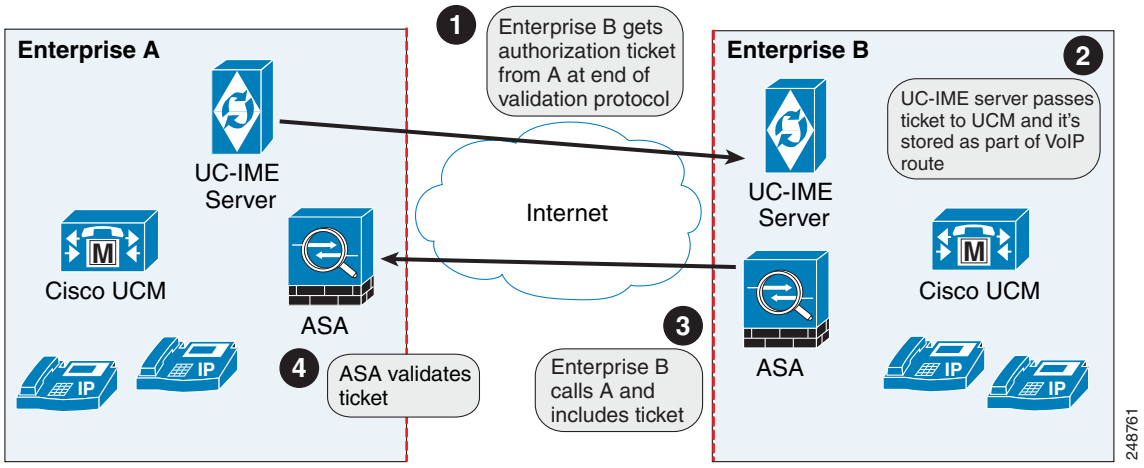

As illustrated in Figure 69-2. Enterprise B makes a PSTN call to enterprise A. That call completes successfully. Later, Enterprise B Cisco Intercompany Media Engine server initiates validation procedures with Enterprise A. These validation procedures succeed. During the validation handshake, Enterprise B sends Enterprise A its domain name. Enterprise A verifies that this domain name is not on the blacklisted set of domains. Assuming it is not, Enterprise A creates a ticket.

Subsequently, someone in Enterprise B calls that number again. That call setup message from Enterprise B to Enterprise A includes the ticket in the X-Cisco-UC-IME-Ticket header field in the SIP INVITE message. This message arrives at the Enterprise A ASA. The ASA verifies the signature and computes several checks on the ticket to make sure it is valid. If the ticket is valid, the ASA forwards the request to Cisco UCM (including the ticket). Because the ASA drops requests that lack a valid ticket, unauthorized calls are never received by Cisco UCM.

The ticket password is a 128 bit random key, which can be thought of as a shared password between the adaptive security appliance and the Cisco Intercompany Media Engine server. This password is generated by the Cisco Intercompany Media Engine server and is used by a Cisco Intercompany Media Engine SIP trunk to generate a ticket to allow a call to be made between Cisco Intercompany Media Engine SIP trunks. A ticket is a signed object that contains a number of fields that grant permission to the calling domain to make a Cisco Intercompany Media Engine call to a specific number. The ticket is signed by the ticket password.

The Cisco Intercompany Media Engine also required that you configure an epoch for the password. The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password.

Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch. If you change the epoch value, the tickets in use at remote enterprises become invalid. The incoming calls from the remote enterprises fallback to the PSTN until the terminating enterprise reissues tickets with the new epoch value and password.

The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. If you change the password or epoch on the ASA, you must update them on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Call Fallback to the PSTN

Cisco Intercompany Media Engine provides features that manage the QoS on the Internet, such as the ability to monitor QoS of the RTP traffic in real-time and fallback to PSTN automatically if problems arise. Call fallback from Internet VoIP calls to the public switched telephone network (PSTN) can occur for two reasons changes in connection quality and signal failure for the Cisco Intercompany Media Engine.

Internet connections can vary wildly in their quality and vary over time. Therefore, even if a call is sent over VoIP because the quality of the connection was good, the connection quality might worsen mid-call. To ensure an overall good experience for the end user, Cisco Intercompany Media Engine attempts to perform a mid-call fallback.

Performing a mid-call fallback requires the adaptive security appliance to monitor the RTP packets coming from the Internet and send information into an RTP Monitoring Algorithm (RMA) API, which will indicates to the adaptive security appliance whether fallback is required. If fallback is required, the adaptive security appliance sends a REFER message to Cisco UCM to tell it that it needs to fallback the call to PSTN.

The TLS signaling connections from the Cisco UCM are terminated on the adaptive security appliance and a TCP or TLS connection is initiated to the Cisco UCM. SRTP (media) sent from external IP phones to the internal network IP phone via the adaptive security appliance is converted to RTP. The adaptive security appliance inserts itself into the media path by modifying the SIP signaling messages that are sent over the SIP trunk between Cisco UCMs. TLS (signaling) and SRTP are always terminated on the adaptive security appliance.

If signaling problems occur, the call falls back to the PSTN; however, the Cisco UCM initiates the PSTN fall back and the adaptive security appliance does not send REFER message.

Architecture and Deployment Scenarios for Cisco Intercompany Media Engine

This section includes the following topics:

- Architecture, page 69-5
- Basic Deployment, page 69-6
- Off Path Deployment, page 69-7

Architecture

Within the enterprise, Cisco Intercompany Media Engine is deployed with the following components for the following purposes:

- The adaptive security appliance—Enabled with the Cisco Intercompany Media Engine Proxy, provides perimeter security functions and inspects SIP signaling between SIP trunks.
- Cisco Intercompany Media Engine (UC-IME) server— Located in the DMZ, provides an automated provisioning service by learning new VoIP routes to particular phone numbers, and recording those routes in Cisco UCM. The Cisco Intercompany Media Engine server does not perform call control.
- Cisco Unified Communications Manager (Cisco UCM)—Responsible for call control and processing. Cisco UCM connects to the Cisco Intercompany Media Engine server by using the Access Protocol to publish and exchange updates. The architecture can consist of a single Cisco UCM or a Cisco UCM cluster within the enterprise.

- Cisco Intercompany Media Engine (UC-IME) Bootstrap server—Provides a certificate required admission onto the public peer-to-peer network for Cisco Intercompany Media Engine.

Figure 69-3 illustrates the components of the Cisco Intercompany Media Engine in a basic deployment.

Figure 69-3 Cisco Intercompany Media Engine Architecture in a Basic Deployment

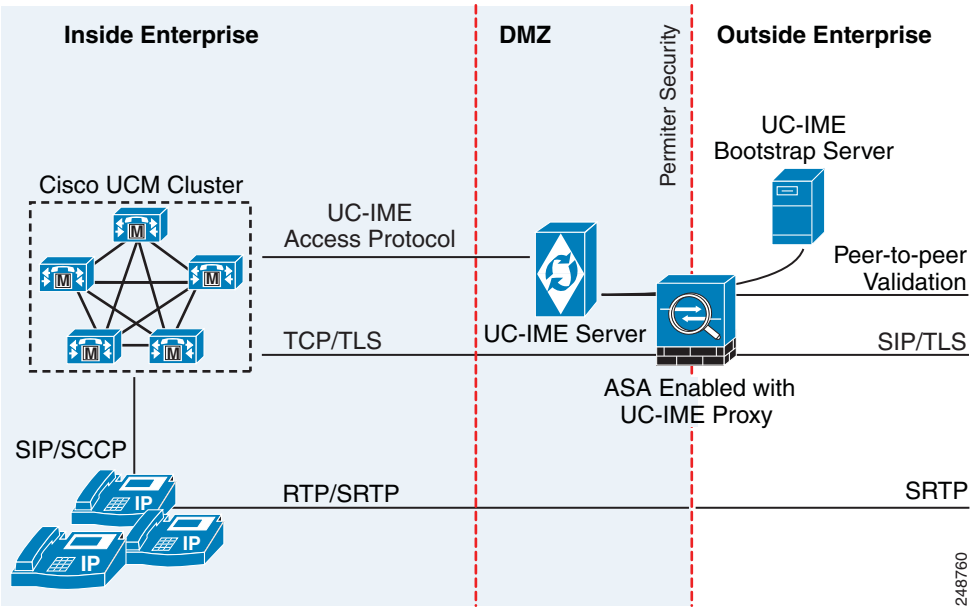

Basic Deployment

In a basic deployment, the Cisco Intercompany Media Engine Proxy sits in-line with the Internet firewall such that all Internet traffic traverses the adaptive security appliance. In this deployment, a single Cisco UCM or a Cisco UCM cluster is centrally deployed within the enterprise, along with a Cisco Intercompany Media Engine server (and perhaps a backup).

As shown in Figure 69-4, the adaptive security appliance sits on the edge of the enterprise and inspects SIP signaling by creating dynamic SIP trunks between enterprises.

Figure 69-4 Basic Deployment Scenario

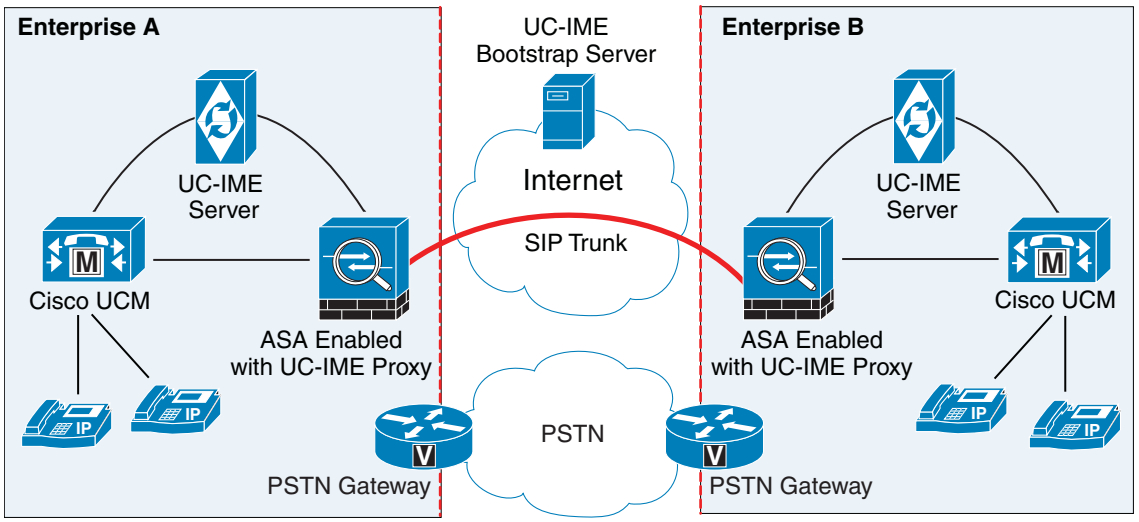

Off Path Deployment

In an off path deployment, inbound and outbound Cisco Intercompany Media Engine calls pass through an adaptive security appliance enabled with the Cisco Intercompany Media Engine Proxy. The adaptive security appliance is located in the DMZ and is configured to support only the Cisco Intercompany Media Engine traffic (SIP signaling and RTP traffic). Normal Internet facing traffic does not flow through this adaptive security appliance.

For all inbound calls, the signaling is directed to the adaptive security appliance because destined Cisco UCMs are configured with the global IP address on the adaptive security appliance. For outbound calls, the called party could be any IP address on the Internet; therefore, the adaptive security appliance is configured with a mapping service that dynamically provides an internal IP address on the adaptive security appliance for each global IP address of the called party on the Internet.

Cisco UCM sends all outbound calls directly to the mapped internal IP address on the adaptive security appliance instead of the global IP address of the called party on the Internet. The adaptive security appliance then forwards the calls to the global IP address of the called party.

Figure 69-5 illustrates the architecture of the Cisco Intercompany Media Engine in an off path deployment.

Figure 69-5 Off Path Deployment of the Adaptive Security Appliance

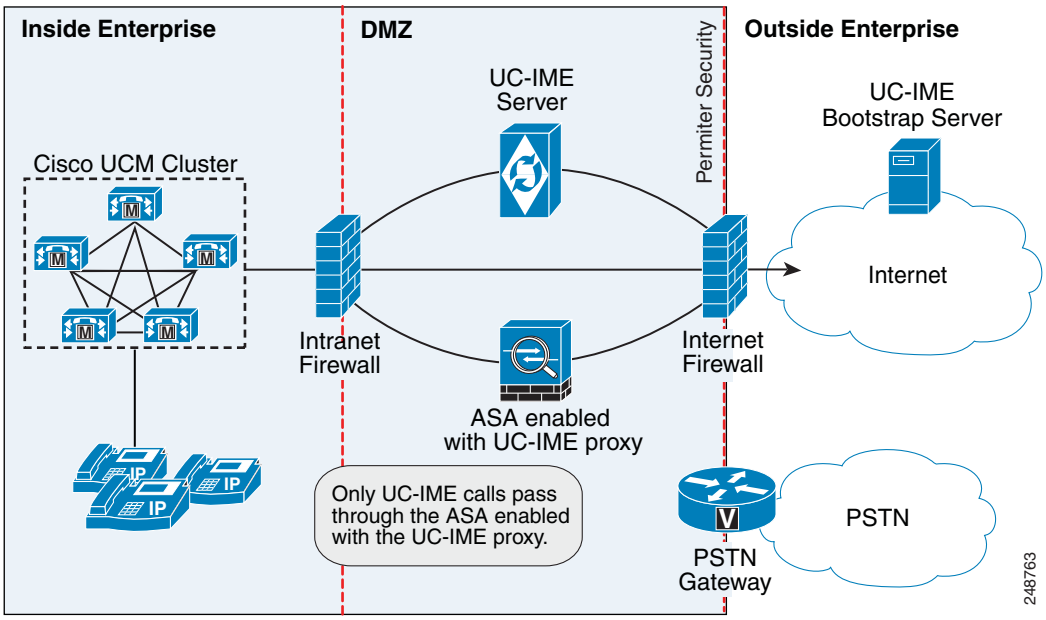

Licensing for Cisco Intercompany Media Engine

The Cisco Intercompany Media Engine feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the details of the Unified Communications Proxy license:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	Intercompany Media Engine license. When you enable the Intercompany Media Engine (IME) license, you can use TLS proxy sessions up to the configured TLS proxy limit. If you also have a Unified Communications (UC) license installed that is higher than the default TLS proxy limit, then the ASA sets the limit to be the UC license limit plus an additional number of sessions depending on your model. You can manually configure the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. If you also install the UC license, then the TLS proxy sessions available for UC are also available for IME sessions. For example, if the configured limit is 1000 TLS proxy sessions, and you purchase a 750-session UC license, then the first 250 IME sessions do not affect the sessions available for UC. If you need more than 250 sessions for IME, then the remaining 750 sessions of the platform limit are used on a first-come, first-served basis by UC and IME. • For a license part number ending in "K8", TLS proxy sessions are limited to 1000. • For a license part number ending in "K9", the TLS proxy limit depends on your configuration and the platform model. **Note** K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted. You might also use SRTP encryption sessions for your connections: • For a K8 license, SRTP sessions are limited to 250. • For a K9 license, there is no limit. **Note** Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

For more information about licensing, see Chapter 5, "Managing Feature Licenses."

Guidelines and Limitations

Context Mode Guidelines

Supported in single context mode only.

Firewall Mode Guidelines

Supported in routed firewall mode only.

IPv6 Guidelines

Does not support IPv6 addresses.

Additional Guidelines and Limitations

Cisco Intercompany Media Engine has the following limitations:

- Fax is not supported. Fax capability needs to be disabled on the SIP trunk.
- Stateful failover of Cisco Unified Intercompany Media Engine is not supported. During failover, existing calls traversing the Cisco Intercompany Media Engine Proxy disconnect; however, new calls successfully traverse the proxy after the failover completes.

- Having Cisco UCMs on more than one of the ASA interfaces is not supported with the Cisco Intercompany Media Engine Proxy. Having the Cisco UCMs on one trusted interface is especially necessary in an off path deployment because the ASA requires that you specify the listening interface for the mapping service and the Cisco UCMs must be connected on one trusted interface.
- Multipart MIME is not supported.
- Only existing SIP features and messages are supported.
- H.264 is not supported.
- RTCP is not supported. The ASA drops any RTCP traffic sent from the inside interface to the outside interface. The ASA does not convert RTCP traffic from the inside interface into SRTP traffic.
- The Cisco Intercompany Media Engine Proxy configured on the ASA creates a dynamic SIP trunk for each connection to a remote enterprise. However, you cannot configure a unique subject name for each SIP trunk. The Cisco Intercompany Media Engine Proxy can have only one subject name configured for the proxy.

  Additionally, the subject DN you configure for the Cisco Intercompany Media Engine Proxy match the domain name that has been set for the local Cisco UCM.
- If a service policy rule for the Cisco Intercompany Media Engine Proxy is removed (by using the no service policy command) and reconfigured, the first call traversing the ASA will fail. The call fails over to the PSTN because the Cisco UCM does not know the connections are cleared and tries to use the recently cleared IME SIP trunk for the signaling.

  To resolve this issue, you must additionally enter the **clear connection all** command and restart the ASA. If the failure is due to failover, the connections from the primary ASA are not synchronized to the standby ASA.
- After the **clear connection all** command is issued on an ASA enabled with a UC-IME Proxy and the IME call fails over to the PSTN, the next IME call between an originating and terminating SCCP IP phone completes but does not have audio and is dropped after the signaling session is established.

  An IME call between SCCP IP phones use the IME SIP trunk in both directions. Namely, the signaling from the calling to called party uses the IME SIP trunk. Then, the called party uses the reverse IME SIP trunk for the return signaling and media exchange. However, this connection is already cleared on the ASA, which causes the IME call to fail.

  The next IME call (the third call after the **clear connection all** command is issued), will be completely successful.

  Note This limitation does not apply when the originating and terminating IP phones are configured with SIP.

- The ASA must be licensed and configured with enough TLS proxy sessions to handle the IME call volume. See Licensing for Cisco Intercompany Media Engine for information about the licensing requirements for TLS proxy sessions.

  This limitation occurs because an IME call cannot fall back to the PSTN when there are not enough TLS proxy sessions left to complete the IME call. An IME call between two SCCP IP phones requires the ASA to use two TLS proxy sessions to successfully complete the TLS handshake.

  Assume for example, the ASA is configured to have a maximum of 100 TLS proxy sessions and IME calls between SCCP IP phones establish 101 TLS proxy sessions. In this example, the next IME call is initiated successfully by the originating SCCP IP phone but fails after the call is accepted by the terminating SCCP IP phone. The terminating IP phone rings and on answering the call, the call hangs due to an incomplete TLS handshake. The call does not fall back to the PSTN.

Configuring Cisco Intercompany Media Engine Proxy

This section contains the following topics:

- Task Flow for Configuring Cisco Intercompany Media Engine, page 69-11
- Configuring NAT for Cisco Intercompany Media Engine Proxy, page 69-12
- Configuring PAT for the Cisco UCM Server, page 69-14
- Creating Access Lists for Cisco Intercompany Media Engine Proxy, page 69-16
- Creating the Media Termination Instance, page 69-17
- Creating the Cisco Intercompany Media Engine Proxy, page 69-18
- Creating Trustpoints and Generating Certificates, page 69-21
- Creating the TLS Proxy, page 69-24
- Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy, page 69-25
- (Optional) Configuring TLS within the Local Enterprise, page 69-27
- (Optional) Configuring Off Path Signaling, page 69-30

Task Flow for Configuring Cisco Intercompany Media Engine

Figure 69-6 provides an example for a basic deployment of the Cisco Intercompany Media Engine. The following tasks include command line examples based on Figure 69-6.

Figure 69-6 Example for Basic (in-line) Deployment Tasks

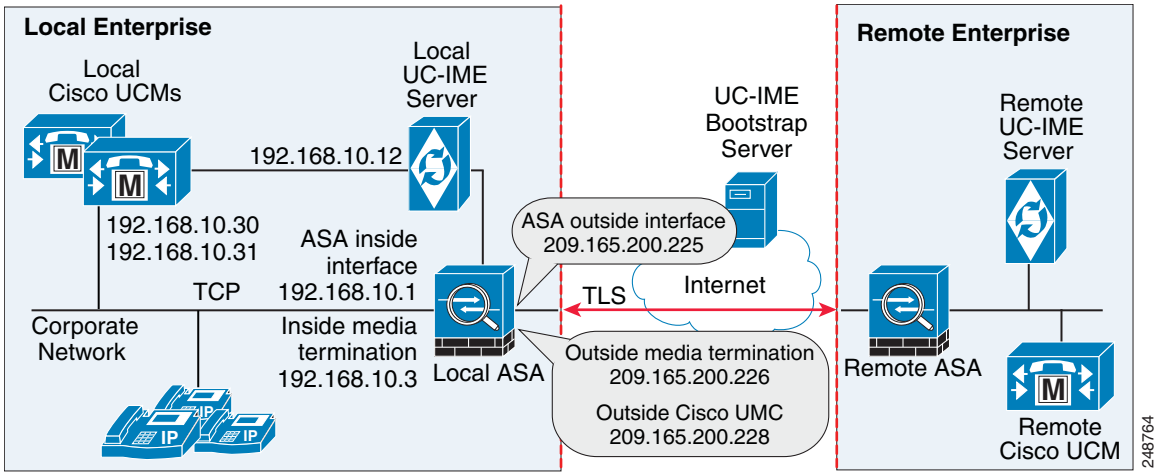

Note Step 1 through Step 8 apply to both basic (in-line) and off path deployments and Step 9 applies only to off path deployment.

To configure a Cisco Intercompany Media Engine for a basic deployment, perform the following tasks.

Step 1 Configure static NAT for Cisco UCM. See Configuring NAT for Cisco Intercompany Media Engine Proxy, page 69-12.

Or

Configure PAT for the UCM server. See Configuring PAT for the Cisco UCM Server, page 69-14.

Step 2 Create access lists for Cisco Intercompany Media Engine Proxy. See Creating Access Lists for Cisco Intercompany Media Engine Proxy, page 69-16.

Step 3 Create the media termination address instance for Cisco Intercompany Media Engine Proxy. See Creating the Media Termination Instance, page 69-17.

Step 4 Create the Cisco Intercompany Media Engine Proxy. See Creating the Cisco Intercompany Media Engine Proxy, page 69-18.

Step 5 Create trustpoints and generate certificates for the Cisco Intercompany Media Engine Proxy. See Creating Trustpoints and Generating Certificates, page 69-21.

Step 6 Create the TLS proxy. See Creating the TLS Proxy, page 69-24.

Step 7 Configure SIP inspection for the Cisco Intercompany Media Engine Proxy. See Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy, page 69-25.

Step 8 (Optional) Configure TLS within the enterprise. See (Optional) Configuring TLS within the Local Enterprise, page 69-27.

Step 9 (Optional) Configure off path signaling. See (Optional) Configuring Off Path Signaling, page 69-30.

Note You only perform Step 9 when you are configuring the Cisco Intercompany Media Engine Proxy in an off path deployment.

Configuring NAT for Cisco Intercompany Media Engine Proxy

To configure auto NAT, you first configure an object; then use the **nat** command in the object configuration mode.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

Alternatively, you can configure PAT for the Cisco Intercompany Media Engine Proxy. See Configuring PAT for the Cisco UCM Server, page 69-14.

Figure 69-7 Example for Configuring NAT for a Deployment

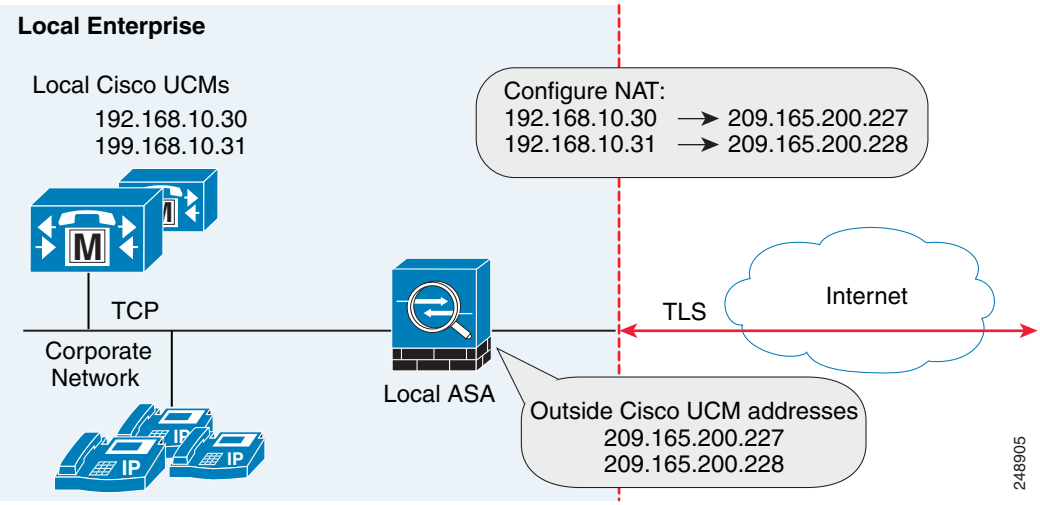

To configure auto NAT rules for the Cisco UCM server, perform the following steps:

	Command	Purpose

Step 1	`hostname(config)# object network name`
	Examples:
	`hostname(config)# object network ucm_real_192.168.10.30`
	`hostname(config)# object network ucm_real_192.168.10.31`
		Configures a network object for the real address of Cisco UCM that you want to translate.

Step 2	`hostname(config-network-object)# host ip_address`
	Examples:
	`hostname(config-network-object)# host 192.168.10.30`
	`hostname(config-network-object)# host 192.168.10.31`
		Specifies the real IP address of the Cisco UCM host for the network object.

Step 3	(Optional)
	`hostname(config-network-object)# description string`
	Example:
	`hostname(config-network-object)# description "Cisco UCM Real Address"`
		Provides a description of the network object.

Step 4	`hostname(config-network-object)# exit`	Exits from the objects configuration mode.

Step 5	`hostname(config)# object network name`
	Example:
	`hostname(config)# object network ucm_map_209.165.200.228`
		Configures a network object for the mapped address of the Cisco UCM.

Step 6	`hostname(config-network-object)# host ip_address`
	Example:
	`hostname(config-network-object)# host 209.165.200.228`
		Specifies the mapped IP address of the Cisco UCM host for the network object.

Step 7	(Optional)
	`hostname(config-network-object)# description string`
	Example:
	`hostname(config-network-object)# description "Cisco UCM Mapped Address"`
		Provides a description of the network object.

	Command	Purpose
Step 8	`hostname(config-network-object)# exit`	Exits from the objects configuration mode.

Step 9	`hostname(config)# nat (inside,outside) source static` *real_obj mapped_obj*
	Examples:
	`hostname(config)# nat (inside,outside) source static ucm_real_192.168.10.30 ucm_209.165.200.228`
	`hostname(config)# nat (inside,outside) source static ucm_real_192.168.10.31 ucm_209.165.200.228`
		Specifies the address translation on the network objects created in this procedure.
		Where *real_obj* is the *name* that you created in Step 1 in this task.
		Where *mapped_obj* is the *name* that you created in Step 5 in this task.

What to Do Next

Create the access lists for the Cisco Intercompany Media Engine Proxy. See Creating Access Lists for Cisco Intercompany Media Engine Proxy, page 69-16.

Configuring PAT for the Cisco UCM Server

Perform this task as an alternative to configuring NAT for the Cisco Intercompany Media Engine Proxy.

Figure 69-8 Example for Configuring PAT for a Deployment

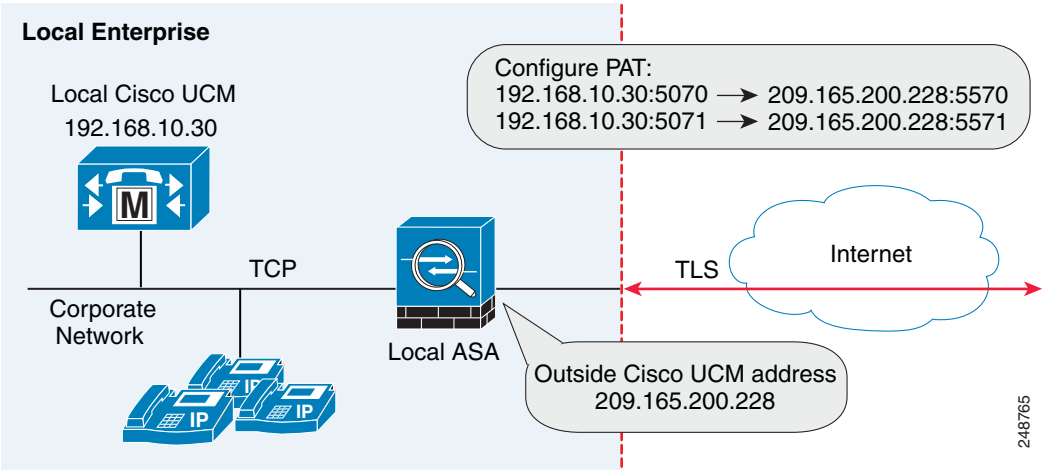

Note You only perform this step when NAT is not configured for the Cisco UCM server.

To configure PAT for the Cisco UCM server, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **object network** *name* **Examples:** hostname(config)# object network ucm-pat-209.165.200.228	Configures a network object for the outside IP address of Cisco UCM that you want to translate.
Step 2	hostname(config-network-object)# **host** *ip_address* **Example:** hostname(config-network-object)# host 209.165.200.228	Specifies the real IP address of the Cisco UCM host for the network object.
Step 3	hostname(config-network-object)# **exit**	Exits from the objects configuration mode.
Step 4	hostname(config)# **object service** *name* **Examples:** hostname(config)# object service tcp_5070 hostname(config)# object service tcp_5071	Creates a service object for the outside Cisco Intercompany Media Engine port.
Step 5	hostname(config-service-object)# **tcp source eq** *port* **Examples:** hostname(config-service-object)# tcp source eq 5070 hostname(config-service-object)# tcp source eq 5071	Specifies the port number.
Step 6	hostname(config-service-object)# **exit**	Exits from the objects configuration mode.
Step 7	hostname(config)# **object network** *name* **Examples:** hostname(config)# object network ucm-real-192.168.10.30 hostname(config)# object network ucm-real-192.168.10.31	Configures a network object to represent the real IP address of Cisco UCM.
Step 8	hostname(config-network-object)# **host** *ip_address* **Examples:** hostname(config-network-object)# host 192.168.10.30 hostname(config-network-object)# host 192.168.10.31	Specifies the real IP address of the Cisco UCM host for the network object.
Step 9	hostname(config-network-object)# **exit**	Exits from the objects configuration mode.
Step 10	hostname(config)# **object service** *name* **Examples:** hostname(config)# object service tcp_5570 hostname(config)# object service tcp_5571	Creates a service objects for Cisco UCM SIP port.
Step 11	hostname(config-service-object)# **tcp source eq** *port* **Example:** hostname(config-service-object)# tcp source eq 5570 hostname(config-service-object)# tcp source eq 5571	Specifies the port number.
Step 12	hostname(config-service-object)# **exit**	Exits from the objects configuration mode.
Step 13	hostname(config)# **nat (inside,outside) source static** *real_obj mapped_obj* **service** *real_port mapped_port* **Examples:** hostname(config)# nat (inside,outside) source static ucm-real-192.168.10.30 ucm-pat-209.165.200.228 service tcp_5070 tcp_5570 hostname(config)# nat (inside,outside) source static ucm-real-192.168.10.31 ucm-pat-128.106.254.5 service tcp_5071 tcp_5571	Creates a static mapping for Cisco UCM. *Where real_obj is the name that you created in Step 1 in this task.* *Where mapped_obj is the name that you created in Step 7 in this task.* *Where real_port is the name that you created in Step 4 in this task.* *Where mapped_obj is the name that you created in Step 10 in this task.*

Creating Access Lists for Cisco Intercompany Media Engine Proxy

To configure access lists for the Cisco Intercompany Media Engine Proxy to reach the Cisco UCM server, perform the following steps.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

	Command	Purpose
Step 1	hostname(config)# **access-list** *id* **extended permit tcp any host** *ip_address* **eq** *port* **Example:** hostname(config)# access-list incoming extended permit tcp any host 192.168.10.30 eq 5070	Adds an Access Control Entry (ACE). An access list is made up of one or more ACEs with the same access list ID. This ACE provides access control by allowing incoming access for Cisco Intercompany Media Engine connections on the specified port. In the *ip_address* argument, provide the real IP address of Cisco UCM.
Step 2	hostname(config)# **access-group** *access-list* **in interface** *interface_name* **Example:** hostname(config)# access-group incoming in interface outside	Binds the access list to an interface.
Step 3	hostname(config)# **access-list** *id* **extended permit tcp any host** *ip_address* **eq** *port* **Example:** hostname(config)# access-list ime-inbound-sip extended permit tcp any host 192.168.10.30 eq 5070	Adds an ACE. This ACE allows the ASA to allow inbound SIP traffic for Cisco Intercompany Media Engine. This entry is used to classify traffic for the class and policy map. **Note** The port that you configure here must match the trunk settings configured on Cisco UCM. See the Cisco Unified Communications Manager documentation for information about this configuration setting.
Step 4	hostname(config)# **access-list** *id* **extended permit tcp** *ip_address mask* **any range** *range* **Example:** hostname(config)# access-list ime-outbound-sip extended permit tcp 192.168.10.30 255.255.255.255 any range 5000 6000	Adds an ACE. This ACE allows the ASA to allow outbound SIP traffic for Cisco Intercompany Media Engine (in the example, any TCP traffic with source as 192.168.10.30 and destination port range between 5000 and 6000). This entry is used to classify traffic for the class and policy map. **Note** Ensure that TCP traffic between Cisco UCM and the Cisco Intercompany Media Engine server does not use this port range (if that connection goes through the ASA).
Step 5	hostname(config)# **access-list** *id* **permit tcp any host** *ip_address* **eq 6084** **Example:** hostname(config)# access-list ime-traffic permit tcp any host 192.168.10.12 eq 6084	Adds an ACE. This ACE allows the ASA to allow traffic from the Cisco Intercompany Media Engine server to remote Cisco Intercompany Media Engine servers.
Step 6	hostname(config)# **access-list** *id* **permit tcp any host** *ip_address* **eq 8470** **Example:** hostname(config)# access-list ime-bootserver-traffic permit tcp any host 192.168.10.12 eq 8470	Adds an ACE. This ACE allows the ASA to allow traffic from the Cisco Intercompany Media Engine server to the Bootstrap server for the Cisco Intercompany Media Engine.

What to Do Next

Create the media termination instance on the ASA for the Cisco Intercompany Media Engine Proxy. See Creating the Media Termination Instance, page 69-17.

Creating the Media Termination Instance

Guidelines

The media termination address you configure must meet these requirements:

- If you decide to configure a media-termination address on interfaces (rather than using a global interface), you must configure a media-termination address on at least two interfaces (the inside and an outside interface) before applying the service policy for the Cisco Intercompany Media Engine Proxy. Otherwise, you will receive an error message when enabling the proxy with SIP inspection.

Note Cisco recommends that you configure the media-termination address for the Cisco Intercompany Media Engine Proxy on interfaces rather than configuring a global media-termination address.

- The Cisco Intercompany Media Engine Proxy can use only one type of media termination instance at a time; for example, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.

Note If you change any Cisco Intercompany Media Engine Proxy settings after you create the media-termination address for the proxy, you must reconfigure the media-termination address by using the **no media-termination** command, and then reconfiguring it as described in this procedure.

Procedure

Create the media termination instance to use with the Cisco Intercompany Media Engine Proxy.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

To create the media termination instance for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **media-termination** *instance_name* **Example:** hostname(config)# **media-termination** *uc-ime-media-term*	Creates the media termination instance that you attach to the Cisco Intercompany Media Engine Proxy.
Step 2	hostname(config-media-termination)# **address** *ip_address* **interface** *intf_name* **Examples:** hostname(config-media-termination)# address 209.165.200.228 interface outside	Configures the media-termination address used by the outside interface of the ASA. The outside IP address must be a publicly routable address that is an unused IP address within the address range on that interface. See Creating the Cisco Intercompany Media Engine Proxy, page 69-18 for information about the UC-IME proxy settings. See CLI configuration guide for information about the **no service-policy** command.
Step 3	hostname(config-media-termination)# **address** *ip_address* **interface** *intf_name* **Examples:** hostname(config-media-termination)# address 192.168.10.3 interface inside	Configures a media termination address used by the inside interface of the ASA. **Note** The IP address must be an unused IP address within the same subnet on that interface.
Step 4	(Optional) hostname(config-media-termination)# ***rtp-min-port*** *port1* ***rtp-maxport*** *port2* **Examples:** hostname(config-media-termination)# rtp-min-port 1000 rtp-maxport 2000	Configures the rtp-min-port and rtp-max-port limits for the Cisco Intercompany Media Engine Proxy. Configure the RTP port range for the media termination point when you need to scale the number of calls that the Cisco Intercompany Media Engine supports. Where *port1* specifies the minimum value for the RTP port range for the media termination point, where port1 can be a value from 1024 to 65535. By default, the value for *port1* is 16384. Where *port2* specifies the maximum value for the RTP port range for the media termination point, where port2 can be a value from 1024 to 65535. By default, the value for *port2* is 32767.

What To Do Next

Once you have created the media termination instance, create the Cisco Intercompany Media Engine Proxy. See Creating the Cisco Intercompany Media Engine Proxy, page 69-18.

Creating the Cisco Intercompany Media Engine Proxy

To create the Cisco Intercompany Media Engine Proxy, perform the following steps.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

Note You cannot change any of the configuration settings for the Cisco Intercompany Media Engine Proxy described in this procedure when the proxy is enabled for SIP inspection. Remove the Cisco Intercompany Media Engine Proxy from SIP inspection before changing any of the settings described in this procedure.

	Command	Purpose

Step 1	`hostname(config)# uc-ime uc_ime_name`
	Example:
	`hostname(config)# uc-ime local-ent-ime`
		Configures the Cisco Intercompany Media Engine Proxy.
		Where *uc_ime_name* is the name of the Cisco Intercompany Media Engine Proxy. The name is limited to 64 characters.
		Only one Cisco Intercompany Media Engine Proxy can be configured on the ASA.

Step 2	`hostname(config-uc-ime)# media-termination mta_instance_name`
	Example:
	`hostname(config-uc-ime)# media-termination ime-media-term`
		Specifies the media termination instance used by the Cisco Intercompany Media Engine Proxy.
		Note You must create the media termination instance before you specify it in the Cisco Intercompany Media Engine Proxy.
		Where *mta_instance_name* is the *instance_name* that you created in Step 1 of Creating the Media Termination Instance.
		See Creating the Media Termination Instance, page 69-17 for the steps to create the media termination instance.

Step 3	`hostname(config-uc-ime)# ucm address ip_address trunk-security-mode [nonsecure | secure]`
	Example:
	`hostname(config-uc-ime)# ucm address 192.168.10.30 trunk-security-mode non-secure`
		Specifies the Cisco UCM server in the enterprise. You must specify the real IP address of the Cisco UCM server. Do not specify a mapped IP address for the server.
		Note You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled.
		Where the **nonsecure** and **secure** options specify the security mode of the Cisco UCM or cluster of Cisco UCMs.
		Note Specifying **secure** for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS; therefore, you must configure TLS for components. See (Optional) Configuring TLS within the Local Enterprise, page 69-27.
		You can specify the **secure** option in this task or you can update it later while configuring TLS for the enterprise. See Step 11 in (Optional) Configuring TLS within the Local Enterprise, page 69-27.

	Command	Purpose
Step 4	hostname(config-uc-ime)# **ticket epoch** *n* **password** *password* **Example:** hostname(config-uc-ime)# ticket epoch 1 password password1234	Configures the ticket epoch and password for Cisco Intercompany Media Engine. Where *n* is an integer from 1-255. The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password. Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch. If you change the epoch value, the current password is invalidated and you must enter a new password. Where *password* contains a minimum of 10 and a maximum of 64 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character. We recommend a password of at least 20 characters. Only one password can be configured at a time. The ticket password is stored onto flash. The output of the **show running-config uc-ime** command displays ***** instead of the password string. **Note** The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

	Command	Purpose

Step 5	(Optional)
	`hostname(config-uc-ime)# fallback monitoring timer timer_millisec | hold-down timer timer_sec`
	Examples:
	`hostname(config-uc-ime)# fallback monitoring timer 120`
	`hostname(config-uc-ime)# fallback hold-down timer 30`
		Specifies the fallback timers for Cisco Intercompany Media Engine.
		Specifying **monitoring timer** sets the time between which the ASA samples the RTP packets received from the Internet. The ASA uses the data sample to determine if fallback to the PSTN is needed for a call.
		Where *timer_millisec* specifies the length of the monitoring timer. By default, the length is 100 milliseconds for the monitoring timer and the allowed range is 10-600 ms.
		Specifying **hold-down timer** sets the amount of time that ASA waits before notifying Cisco UCM whether to fall back to PSTN.
		Where *timer_sec* specifies the length of the hold-down timer. By default, the length is 20 seconds for the hold-down timer and the allowed range is 10-360 seconds.
		If you do not use this command to specify fallback timers, the ASA uses the default settings for the fallback timers.

Step 6	(Optional)
	`hostname(config-uc-ime)# fallback sensitivity-file file_name`
	Example:
	`hostname(config-uc-ime)# fallback sensitivity-file ime-fallback-sensitvity.fbs`
		Specifies the file to use for mid-call PSTN fallback.
		Where *file_name* must be the name of a file on disk that includes the .fbs file extension.
		The fallback file is used to determine whether the QoS of the call is poor enough for the Cisco Intercompany Media Engine to move the call to the PSTN.

What to Do Next

Install the certificate on the local entity truststore. You could also enroll the certificate with a local CA trusted by the local entity.

Creating Trustpoints and Generating Certificates

You need to generate the keypair for the certificate used by the ASA, and configure a trustpoint to identify the certificate sent by the ASA in the TLS handshake.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

Note This task instructs you on how to create trustpoints for the local enterprise and the remote enterprise and how to exchange certificates between these two enterprises. This task does not provide steps for creating trustpoints and exchanging certificates between the local Cisco UCM and the local ASA. However, if you require additional security within the local enterprise, you must perform the optional task (Optional)

Configuring TLS within the Local Enterprise, page 69-27. Performing that task allows for secure TLS connections between the local Cisco UCM and the local ASA. The instructions in that task describe how to create trustpoints between the local Cisco UCM and the local ASA.

Prerequisites for Installing Certificates

To create a proxy certificate on the ASA that is trusted by the remote entity, obtain a certificate from a trusted CA or export it from the remote enterprise ASA.

To export the certificate from the remote enterprise, you enter the following command on the remote ASA:

hostname(config)# **crypto ca export trustpoint identity-certificate**

The ASA prompts displays the certificate in the terminal screen. Copy the certificate from the terminal screen. You will need the certificate text in Step 5 of this task.

Procedure

To create the trustpoints and generate certificates, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **crypto key generate rsa label** *key-pair-label* **modulus** *size* **Example:** hostname(config)# crypto key generate rsa label local-ent-key modulus 2048	On the local ASA, creates the RSA keypair that can be used for the trustpoints. This is the keypair and trustpoint for the local entities signed certificate. The modulus key size that you select depends on the level of security that you want to configure and on any limitations imposed by the CA from which you are obtaining the certificate. The larger the number that you select, the higher the security level will be for the certificate. Most CAs recommend 2048 for the key modulus size; however, **Note** GoDaddy requires a key modulus size of 2048.
Step 2	hostname(config)# **crypto ca trustpoint** *trustpoint_name* **Example:** hostname(config)# crypto ca trustpoint local_ent	Enters the trustpoint configuration mode for the specified trustpoint so that you can create the trustpoint for the local entity. A trustpoint represents a CA identity and possibly a device identity, based on a certificate issued by the CA. Maximum name length is 128 characters.
Step 3	hostname(config-ca-trustpoint)# **subject-name** *X.500_name* **Example:** hostname(config-ca-trustpoint)# subject-name cn=Ent-local-domain-name**	Includes the indicated subject DN in the certificate during enrollment. **Note** The domain name that you enter here must match the domain name that has been set for the local Cisco UCM. For information about how to configure the domain name for Cisco UCM, see the Cisco Unified Communications Manager documentation for information.

	Command	Purpose

Step 4	`hostname(config-ca-trustpoint)# keypair keyname`
	Example:
	`hostname(config-ca-trustpoint)# keypair local-ent-key`
		Specifies the key pair whose public key is to be certified.

Step 5	`hostname(config-ca-trustpoint)# enroll terminal`	Specifies that you will use the "copy and paste" method of enrollment with this trustpoint (also known as manual enrollment).
Step 6	`hostname(config-ca-trustpoint)# exit`	Exits from the CA Trustpoint configuration mode.

Step 7	`hostname(config)# crypto ca enroll trustpoint`
	Example:
	`hostname(config)# crypto ca enroll remote-ent`
	`%`
	`% Start certificate enrollment ...`
	`% The subject name in the certificate will be:`
	`% cn=enterpriseA`
	`% The fully-qualified domain name in the certificate will`
	`@ be: ciscoasa`
	`% Include the device serial number in the subject name?`
	`[yes/no]: no`
	`Display Certificate Request to terminal? [yes/no]: yes`
		Starts the enrollment process with the CA.
		Where *trustpoint* is the same as the value you entered for *trustpoint_name* in Step 2.
		When the trustpoint is configured for manual enrollment (**enroll terminal** command), the ASA writes a base-64-encoded PKCS10 certification request to the console and then displays the CLI prompt. Copy the text from the prompt.
		Submit the certificate request to the CA, for example, by pasting the text displayed at the prompt into the certificate signing request enrollment page on the CA website.
		When the CA returns the signed identity certificate, proceed to Step 8 in this procedure.

Step 8	`hostname(config)# crypto ca import trustpoint certificate`
	Example:
	`hostname(config)# crypto ca import remote-ent certificate`
		Imports the signed certificate received from the CA in response to a manual enrollment request.
		Where *trustpoint* specifies the trustpoint you created in Step 2.
		The ASA prompts you to paste the base-64 formatted signed certificate onto the terminal.

Step 9	`hostname(config)# crypto ca authenticate trustpoint`
	Example:
	`hostname(config)# crypto ca authenticate remote-ent`
		Authenticates the third-party identity certificate received from the CA. The identity certificate is associated with a trustpoint created for the remote enterprise.
		The ASA prompts you to paste the base-64 formatted identity certificate from the CA onto the terminal.

What to Do Next

Create the TLS proxy for the Cisco Intercompany Media Engine. See the "Creating the TLS Proxy" section on page 69-24.

Creating the TLS Proxy

Because either enterprise, namely the local or remote Cisco UCM servers, can initiate the TLS handshake (unlike IP Telephony or Cisco Mobility Advantage, where only the clients initiate the TLS handshake), you must configure by-directional TLS proxy rules. Each enterprise can have an ASA as the TLS proxy.

Create TLS proxy instances for the local and remote entity initiated connections respectively. The entity that initiates the TLS connection is in the role of "TLS client." Because the TLS proxy has a strict definition of "client" and "server" proxy, two TLS proxy instances must be defined if either of the entities could initiate the connection.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

To create the TLS proxy, perform the following steps:

	Command	Purpose

Step 1	`hostname(config)# tls-proxy proxy_name`
	Example:
	`hostname(config)# tls-proxy local_to_remote-ent`
		Creates the TLS proxy for the outbound connections.

Step 2	`hostname(config-tlsp)# client trust-point proxy_trustpoint`
	Example:
	`hostname(config-tlsp)# client trust-point local-ent`
		For **outbound** connections, specifies the trustpoint and associated certificate that the adaptive security appliance uses in the TLS handshake when the adaptive security appliance assumes the role of the TLS client. The certificate must be owned by the adaptive security appliance (identity certificate).
		Where *proxy_trustpoint* specifies the trustpoint defined by the **crypto ca trustpoint** command in Step 2 in "Creating Trustpoints and Generating Certificates" section on page 69-21.

Step 3	`hostname(config-tlsp)# client cipher-suite cipher_suite`
	Example:
	`hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1`
		For outbound connections, controls the TLS handshake parameter for the cipher suite.
		Where `cipher_suite` includes des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, or null-sha1.
		For client proxy (the proxy acts as a TLS client to the server), the user-defined cipher suite replaces the default cipher suite, or the one defined by the **ssl encryption** command. Use this command to achieve difference ciphers between the two TLS sessions. You should use AES ciphers with the Cisco UCM server.

Step 4	`hostname(config-tlsp)# exit`	Exits from the TLS proxy configuration mode.

Step 5	`hostname(config)# tls-proxy proxy_name`
	Example:
	`hostname(config)# tls-proxy remote_to_local-ent`
		Create the TLS proxy for inbound connections.

	Command	Purpose

Step 6	`hostname(config-tlsp)# server trust-point proxy_trustpoint`
	Example:
	`hostname(config-tlsp)# server trust-point local-ent`
		For **inbound** connections, specifies the proxy trustpoint certificate presented during TLS handshake. The certificate must be owned by the adaptive security appliance (identity certificate).
		Where *proxy_trustpoint* specifies the trustpoint defined by the **crypto ca trustpoint** command in Step 2 in "Creating Trustpoints and Generating Certificates" section on page 69-21.
		Because the TLS proxy has strict definition of client proxy and server proxy, two TLS proxy instances must be defined if either of the entities could initiate the connection.

Step 7	`hostname(config-tlsp)# client cipher-suite cipher_suite`
	Example:
	`hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1`
		For inbound connections, controls the TLS handshake parameter for the cipher suite.
		Where *cipher_suite* includes des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, or null-sha1.

Step 8	`hostname(config-tlsp)# exit`	Exits from the TSL proxy configuration mode.
Step 9	`hostname(config)# ssl encryption 3des-sha1 aes128-sha1 [algorithms]`	Specifies the encryption algorithms that the SSL/TLS protocol uses. Specifying the 3des-shal and aes128-shal is required. Specifying other algorithms is optional. **Note** The Cisco Intercompany Media Engine Proxy requires that you use strong encryption. You must specify this command when the proxy is licensed using a K9 license.

What to Do Next

Once you have created the TLS proxy, enable it for SIP inspection.

Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy

Enable the TLS proxy for SIP inspection and define policies for both entities that could initiate the connection.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 69-6 on page 69-11 for an illustration explaining the example command lines in this task.

Note If you want to change any Cisco Intercompany Media Engine Proxy settings after you enable SIP inspection, you must enter the **no service-policy** command, and then reconfigure the service policy as described in this procedure. Removing and reconfiguring the service policy does not affect existing calls; however, the first call traversing the Cisco Intercompany Media Engine Proxy will fail. Enter the **clear connection** command and restart the ASA.

To enable SIP inspection for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **class-map** *class_map_name* **Examples:** hostname(config)# class-map ime-inbound-sip	Defines a class for the inbound Cisco Intercompany Media Engine SIP traffic.
Step 2	hostname(config-cmap)# **match access-list** *access_list_name* **Examples:** hostname(config-cmap)# match access-list ime-inbound-sip	Identifies the SIP traffic to inspect. Where the *access_list_name* is the access list you created in Step 3, page 69-16 of the task Creating Access Lists for Cisco Intercompany Media Engine Proxy.
Step 3	hostname(config-cmap)# **exit**	Exits from the class map configuration mode.
Step 4	hostname(config)# **class-map** *class_map_name* **Examples:** hostname(config)# class-map ime-outbound-sip	Defines a class for the outbound SIP traffic from Cisco Intercompany Media Engine.
Step 5	hostname(config)# **match access-list** *access_list_name* **Examples:** hostname(config-cmap)# match access-list ime-outbound-sip	Identifies which outbound SIP traffic to inspect. Where the *access_list_name* is the access list you created in Step 4, page 69-16 of the task Creating Access Lists for Cisco Intercompany Media Engine Proxy.
Step 6	hostname(config-cmap)# **exit**	Exits from the class map configuration mode.
Step 7	hostname(config)# **policy-map** *name* **Examples:** hostname(config)# policy-map ime-policy	Defines the policy map to which to attach the actions for the class of traffic.
Step 8	hostname(config-pmap)# **class** *classmap_name* **Examples:** hostname(config-pmap)# class ime-outbound-sip	Assigns a class map to the policy map so that you can assign actions to the class map traffic. Where *classmap_name* is the name of the SIP class map that you created in Step 1 in this task.
Step 9	hostname(config-pmap-c)# **inspect sip** [*sip_map*] **tls-proxy** *proxy_name* **uc-ime** *uc_ime_map* **Examples:** hostname(config-pmap-c)# inspect sip tls-proxy local_to_remote-ent uc-ime local-ent-ime	Enables the TLS proxy and Cisco Intercompany Media Engine Proxy for the specified SIP inspection session.
Step 10	hostname(config-cmap-c)# **exit**	Exits from the policy map class configuration mode.
Step 11	hostname(config-pmap)# **class** *class_map_name* **Examples:** hostname(config-pmap)# class ime-inbound-sip	Assigns a class map to the policy map so that you can assign actions to the class map traffic. Where *classmap_name* is the name of the SIP class map that you created in Step 4 in this task.
Step 12	hostname(config-pmap-c)# **inspect sip** [*sip_map*] **tls-proxy** *proxy_name* **uc-ime** *uc_ime_map* **Examples:** hostname(config-pmap-c)# inspect sip tls-proxy remote-to-local-ent uc-ime local-ent-ime	Enables the TLS proxy and Cisco Intercompany Media Engine Proxy for the specified SIP inspection session.
Step 13	hostname(config-pmap-c)# **exit**	Exits from the policy map class configuration mode.

	Command	Purpose
Step 14	`hostname(config-pmap)# exit`	Exits from the policy map configuration mode.

Step 15	`hostname(config)# service-policy policymap_name global`
	Examples:
	`hostname(config)# service-policy ime-policy global`
		Enables the service policy for SIP inspection for all interfaces.
		Where *policymap_name* is the name of the policy map you created in Step 7 of this task.
		See Creating the Cisco Intercompany Media Engine Proxy, page 69-18 for information about the UC-IME proxy settings. See CLI configuration guide for information about the **no service-policy** command.

What to Do Next

Once you have enabled the TLS proxy for SIP inspection, if necessary, configure TLS within the enterprise. See (Optional) Configuring TLS within the Local Enterprise, page 69-27.

(Optional) Configuring TLS within the Local Enterprise

This task is not required if TCP is allowable within the inside network.

TLS within the enterprise refers to the security status of the Cisco Intercompany Media Engine trunk as seen by the ASA.

Note If the transport security for the Cisco Intercompany Media Engine trunk changes on Cisco UCM, it must be changed on the ASA as well. A mismatch will result in call failure. The ASA does not support SRTP with non-secure IME trunks. The ASA assumes SRTP is allowed with secure trunks. So 'SRTP Allowed' must be checked for IME trunks if TLS is used. The ASA supports SRTP fallback to RTP for secure IME trunk calls.

Prerequisites

On the local Cisco UCM, download the Cisco UCM certificate. See the Cisco Unified Communications Manager documentation for information. You will need this certificate when performing Step 6 of this procedure.

Procedure

To configure TLS within the local enterprise, perform the following steps on the local ASA:

	Commands	Purpose
Step 1	hostname(config)# **crypto key generate rsa label** *key-pair-label* hostname(config)# **crypto ca trustpoint** *trustpoint_name* hostname(config-ca-trustpoint)# **enroll self** hostname(config-ca-trustpoint)# **keypair** *keyname* hostname(config-ca-trustpoint)# **subject-name** *x.500_name* **Example:** hostname(config)# crypto key generate rsa label local-ent-key hostname(config)# crypto ca trustpoint local-asa hostname(config-ca-trustpoint)# enroll self hostname(config-ca-trustpoint)# keypair key-local-asa hostname(config-ca-trustpoint)# subject-name *cn=Ent-local-domain-name***., o="Example Corp"	Creates an RSA key and trustpoint for the self-signed certificate. Where *key-pair-label* is the RSA key for the local ASA. Where *trustpoint_name* is the trustpoint for the local ASA. Where *keyname* is key pair for the local ASA. Where *x.500_name* includes the X.500 distinguished name of the local ASA; for example, *cn=Ent-local-domain-name***. **Note** The domain name that you enter here must match the domain name that has been set for the local Cisco UCM. For information about how to configure the domain name for Cisco UCM, see the Cisco Unified Communications Manager documentation for information.
Step 2	hostname(config-ca-trustpoint)# **exit**	Exits from Trustpoint Configuration mode.
Step 3	hostname(config)# **crypto ca export** *trustpoint* **identity-certificate** **Example:** hostname(config)# crypto ca export local-asa identity-certificate	Exports the certificate you created in Step 1. The certificate contents appear on the terminal screen. Copy the certificate from the terminal screen. This certificate enables Cisco UCM to validate the certificate that the ASA sends in the TLS handshake. On the local Cisco UCM, upload the certificate into the Cisco UCM trust store. See the Cisco Unified Communications Manager documentation for information. **Note** The subject name you enter while uploading the certificate to the local Cisco UCM is compared with the X.509 Subject Name field entered on the SIP Trunk Security Profile on Cisco UCM. For example, "Ent-local-domain-name" was entered in Step 1 of this task; therefore, "Ent-local-domain-name" should be entered in the Cisco UCM configuration.
Step 4	hostname(config)# **crypto ca trustpoint** *trustpoint_name* hostname(config-ca-trustpoint)# **enroll terminal** **Example:** hostname(config)# crypto ca trustpoint local-ent-ucm hostname(config-ca-trustpoint)# enroll terminal	Creates a trustpoint for local Cisco UCM. Where *trustpoint_name* is the trustpoint for the local Cisco UCM.
Step 5	hostname(config-ca-trustpoint)# **exit**	Exits from Trustpoint Configuration mode.

	Commands	Purpose
Step 6	hostname(config)# **crypto ca authenticate** *trustpoint* **Example:** hostname(config)# crypto ca authenticate local-ent-ucm	Imports the certificate from local Cisco UCM. Where *trustpoint* is the trustpoint for the local Cisco UCM. Paste the certificate downloaded from the local Cisco UCM. This certificate enables the ASA to validate the certificate that Cisco UCM sends in the TLS handshake.
Step 7	hostname(config)# **tls-proxy** *proxy_name* hostname(config-tlsp)# **server trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client cipher-suite** aes128-sha1 aes256-sha1 3des-sha1 null-sha1 **Example:** hostname(config)# tls-proxy local_to_remote-ent hostname(config-tlsp)# server trust-point local-ent-ucm hostname(config-tlsp)# client trust-point local-ent hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	Updates the TLS proxy for **outbound** connections. Where *proxy_name* is the name you entered in Step 1 of the task Creating the TLS Proxy. Where *proxy_trustpoint* for the **server trust-point** command is the name you entered in Step 4 of this procedure. Where *proxy_trustpoint* for the **client trust-point** command is the name you entered in Step 2 of the task Creating Trustpoints and Generating Certificates. **Note** In this step, you are creating different trustpoints for the client and the server.
Step 8	hostname(config-tlsp)# **exit**	Exits from TLS Proxy Configuration mode.
Step 9	hostname(config)# **tls-proxy** proxy_name hostname(config-tlsp)# **server trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client cipher-suite** aes128-sha1 aes256-sha1 3des-sha1 null-sha1 **Example:** hostname(config)# tls-proxy remote_to_local-ent hostname(config-tlsp)# server trust-point local-ent hostname(config-tlsp)# client trust-point local-ent-ucm hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	Updates the TLS proxy for **inbound** connections. Where *proxy_name* is the name you entered in Step 5 of the task Creating the TLS Proxy. Where *proxy_trustpoint* for the **server trust-point** command is the name you entered in Step 2 of the task Creating Trustpoints and Generating Certificates. Where *proxy_trustpoint* for the **client trust-point** command is the name you entered in Step 4 of this procedure.
Step 10	hostname(config-tlsp)# **exit**	Exits from TLS Proxy Configuration mode.
Step 11	hostname(config)# **uc-ime** *uc_ime_name* hostname(config-uc-ime)# **ucm address** *ip_address* **trunk-security-mode secure** **Example:** hostname(config)# uc-ime local-ent-ime hostname(config-uc-ime)# ucm address 192.168.10.30 trunk-security-mode secure	Updates the Cisco Intercompany Media Engine Proxy for trunk-security-mode. Where *uc_ime_name* is the name you entered in Step 1 of the task Creating the Cisco Intercompany Media Engine Proxy. Only perform this step if you entered nonsecure in Step 3 of the task Creating the Cisco Intercompany Media Engine Proxy.

What to Do Next

Once you have configured the TLS within the enterprise, if necessary, configure off path signaling for an off path deployment. See (Optional) Configuring Off Path Signaling, page 69-30.

(Optional) Configuring Off Path Signaling

Perform this task only when you are configuring the Cisco Intercompany Media Engine Proxy as part of an off path deployment. You might choose to have an off path deployment when you want to use the Cisco Intercompany Media Engine but do not want to replace your existing Internet firewall with an ASA enabled with the Cisco Intercompany Media Engine Proxy.

In an off path deployment, the existing firewall that you have deployed in your environment is not capable of transmitting Cisco Intercompany Media Engine traffic.

Off path signaling requires that outside IP addresses translate to an inside IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address; therefore, using the dynamic NAT configuration on outbound calls, Cisco UCM sends SIP traffic to this internal IP address, and the ASA uses that mapping to determine the real destination on inbound calls. The static NAT or PAT mapping is used for inbound calls in an off path configuration.

Figure 69-9 Example for Configuring Off Path Signaling in an Off Path Deployment

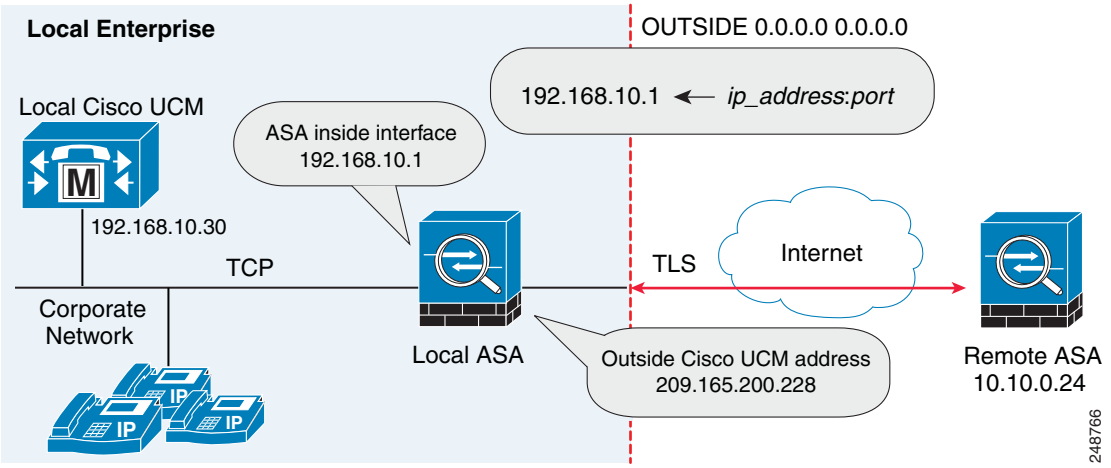

After you configure off path signaling, the ASA mapping service listens on interface “inside” for requests. When it receives a request, it creates a dynamic mapping for the “outside” as the destination interface.

To configure off path signaling for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	`hostname(config)# object network` *name* **Example:** `hostname(config)# object network outside-any`	For the off path ASA, creates a network object to represent all outside addresses.
Step 2	`hostname(config-network-object)# subnet` *ip_address* **Example:** `hostname(config-network-object)# subnet 0.0.0.0 0.0.0.0`	Specifies the IP address of the subnet.
Step 3	`hostname(config-network-object)# nat (outside,inside) dynamic interface inside`	Creates a mapping for the Cisco UCM of remote enterprises.
Step 4	`hostname(config-network-object)# exit`	Exits from the objects configuration mode.

	Command	Purpose

Step 5	`hostname(config)# uc-ime uc_ime_name`
	Example:
	`hostname(config)# uc-ime local-ent-ime`
		Specifies the Cisco Intercompany Media Engine Proxy that you created in the task Creating the Cisco Intercompany Media Engine Proxy, page 69-18.
		Where *uc_ime_name* is the name you specified in Step 1 of Creating the Cisco Intercompany Media Engine Proxy, page 69-18.

Step 6	`hostname(config)# mapping-service listening-interface interface_name [listening-port port] uc-ime-interface uc-ime-interface_name`
	Example:
	`hostname(config-uc-ime)# mapping-service listening-interface inside listening-port 8060 uc-ime-interface outside`
		For the off path ASA, adds the mapping service to the Cisco Intercompany Media Engine Proxy.
		Specifies the interface and listening port for the ASA mapping service.
		You can only configure one mapping server for the Cisco Intercompany Media Engine Proxy.
		Where *interface_name* is the name of the interface on which the ASA listens for the mapping requests.
		Where port is the TCP port on which the ASA listens for the mapping requests. The port number must be between 1024 and 65535 to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.
		Where *uc-ime-interface_name* is the name of the interface that connects to the remote Cisco UCM.

This section contains the following sections:

- Configuring the Cisco UC-IMC Proxy by using the UC-IME Proxy Pane, page 69-31
- Configuring the Cisco UC-IMC Proxy by using the Unified Communications Wizard, page 69-33

Configuring the Cisco UC-IMC Proxy by using the UC-IME Proxy Pane

Use the Configure Cisco Intercompany Media Engine (UC-IME) proxy pane to add or edit a Cisco Intercompany Media Engine Proxy instance.

Note The Cisco Intercompany Media Engine Proxy does not appear as an option under the Unified Communications section of the navigation pane unless the license required for this proxy is installed on the ASA.

Use this pane to create the proxy instance; however, for the UC-IME proxy to be fully functionally, you must complete additional tasks, such as create the required NAT statements, access lists, and MTA, set up the certificates, create the TLS Proxy, and enable SIP inspection.

Depending on whether the UC-IME proxy is deployed off path or in-line of Internet traffic, you must create the appropriate network objects with embedded NAT/PAT statements for the Cisco UCMs.

This pane is available from the Configuration > Firewall > Unified Communications > UC-IME Proxy.

Step 1 Open the Configuration > Firewall > Unified Communications > UC-IME Proxy pane.

Step 2 Check the Enable Cisco UC-IME proxy check box to enable the feature.

Step 3 In the Unified CM Servers area, enter an IP address or hostname for the Cisco Unified Communications Manager (Cisco UCM) or click the ellipsis to open a dialog and browse for an IP address or hostname.

Step 4 In the Trunk Security Mode field, click a security option. Specifying **secure** for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS.

Step 5 Click **Add** to add the Cisco UCM for the Cisco Intercompany Media Engine Proxy. You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled.

Step 6 In the Ticket Epoch field, enter an integer from 1-255.

The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password.

Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch.

If you change the epoch value, the current password is invalidated and you must enter a new password.

Note The epoch and password that you configure in this step on the ASA must match the epoch and password that you configure on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Step 7 In the Ticket Password field, enter a minimum of 10 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character. The ticket password can be up to 64 characters. Confirm the password you entered. Only one password can be configured at a time.

Step 8 Check the Apply MTA to UC-IME Link proxy check box to associate the media termination address with the Cisco Intercompany Media Engine Proxy.

Note You must create the media termination instance before you associate it with the Cisco Intercompany Media Engine Proxy. If necessary, click the Configure MTA button to configure a media termination address instance.

Step 9 If the Cisco Intercompany Media Engine Proxy is being configured as part of off path deployment, check the Enable off path address mapping service checkbox and configure the off path deployment settings:

a. From the Listening Interface field, select an ASA interface. This is the interface on which the ASA listens for the mapping requests.

b. In the Port field, enter a number between 1024 and 65535 as the TCP port on which the ASA listens for the mapping requests. The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

c. From the UC-IME Interface field, select an interface from the list. This is the interface that the ASA uses to connect to the remote Cisco UCM.

Note In an off path deployment any existing ASA that you have deployed in your environment are not capable of transmitting Cisco Intercompany Media Engine traffic. Off-path signaling requires that outside addresses are translated (using NAT) to an inside IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address.

Step 10 In the Fallback area, configure the fallback timer for the Cisco Intercompany Media Engine by specifying the following settings:

a. In the Fallback Sensitivity File field, enter the path to a file in flash memory that the ASA uses for mid-call PSTN fallback. The file name that you enter must be the name of a file on disk that includes the .fbs file extension. Alternatively, click the Browse Flash button to locate and select the file from flash memory.

b. In the Call Quality Evaluation Interval field, enter a number between 10-600 (in milliseconds). This number controls the frequency at which the ASA samples the RTP packets received from the Internet. The ASA uses the data sample to determine if fallback to the PSTN is needed for a call. By default, the length is 100 milliseconds for the timer.

c. In the Notification Interval field, enter a number between 10-360 (in seconds). This number controls the amount of time that the ASA waits before notifying Cisco UCM whether to fall back to PSTN. By default, the length is 20 seconds for this timer.

Note When you change the fallback timer for the Cisco Intercompany Media Engine Proxy, ASDM automatically removes the proxy from SIP inspection and then reapplies SIP inspection when the proxy is re-enabled.

Step 11 Click Apply to save the configuration changes for the Cisco Intercompany Media Engine Proxy.

Configuring the Cisco UC-IMC Proxy by using the Unified Communications Wizard

To configure the Cisco Intercompany Media Engine Proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Intercompany Media Engine Proxy option under the Business-to-Business section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Intercompany Media Engine proxy, importing and installing the required certificates, and finally enables the SIP inspection for the Intercompany Media Engine traffic automatically.

The wizard guides you through these steps to create the Cisco Intercompany Media Engine Proxy:

Step 1 Select the Intercompany Media Engine Proxy option.

Step 2 Select the topology of the Cisco Intercompany Media Engine Proxy, namely whether the ASA is an edge firewall with all Internet traffic flowing through it or whether the ASA is off the path of the main Internet traffic (referred to as an off path deployment).

Step 3 Specify private network settings such as the Cisco UCM IP addresses and the ticket settings.

Step 4 Specify the public network settings.

Step 5 Specify the media termination address settings of Cisco UCM.

Step 6 Configure the local-side certificate management, namely the certificates that are exchanged between the local Cisco Unified Communications Manager servers and the ASA. The identity certificate that the wizard generates in this step needs to be installed on each Cisco Unified Communications Manager (UCM) server in the cluster with the proxy and each identity certificate from the Cisco UCMs need to be installed on the ASA. The certificates are used by the ASA and the Cisco UCMs to authenticate each other, respectively, during TLS handshakes. The wizard only supports self-signed certificates for this step.

Step 7 Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. In this step, the wizard generates a certificate signing request (CSR). After successfully generating the identity certificate request for the proxy, the wizard prompts you to save the file.

You must send the CSR text file to a certificate authority (CA), for example, by pasting the text file into the CSR enrollment page on the CA website. When the CA returns the Identity Certificate, you must install it on the ASA. This certificate is presented to remote servers so that they can authenticate the ASA as a trusted server.

Finally, this step of the wizard assists you in installing the root certificates of the CA from the remote servers so that the ASA can determine that the remote servers are trusted.

The wizard completes by displaying a summary of the configuration created for Cisco Intercompany Media Engine. See the Unified Communications Wizard section in this documentation for more information.

This section describes how to certain options of the **show uc-ime** command to obtain troubleshooting information for the Cisco Intercompany Media Engine Proxy. See the command reference for detailed information about the syntax for these commands.

show uc-ime signaling-sessions

Displays the corresponding SIP signaling sessions stored by the Cisco Intercompany Media Engine Proxy. Use this command to troubleshoot media or signaling failure. The command also displays the fallback parameters extracted from the SIP message headers, whether RTP monitoring is enabled or disabled, and whether SRTP keys are set.

Through the use of the Cisco Intercompany Media Engine Proxy, not only signaling but also media is secured for communication. It provides signaling encryption and SRTP/RTP conversion with SRTP enforced on the Internet side. The Cisco Intercompany Media Engine Proxy inserts itself into the media path by modifying the SIP signaling messages from Cisco UCMs.The Cisco Intercompany Media Engine Proxy sits on the edge of the enterprise and inspects SIP signaling between SIP trunks created between enterprises. It terminates TLS signaling from the Internet and initiates TCP or TLS to the local Cisco UCM.

```
hostname# show uc-ime signaling-sessions
 1 in use, 3 most used
 inside 192.168.10.30:39608 outside 10.194.108.118:5070
   Local Media (audio) conn: 10.194.108.119/29824 to 10.194.108.109/21558
     Local SRTP key set : Remote SRTP key set
   Remote Media (audio) conn: 192.168.10.51/19520 to 192.168.10.3/30930
   Call-ID: ab6d7980-a7d11b08-50-1e0aa8c0@192.168.10.30
   FB Sensitivity: 3
   Session ID: 2948-32325449-0@81a985c9-f3a1-55a0-3b19-96549a027259
```

```
SIP Trunk URI: 81a985c9-f3a1-55a0-3b19-9654@UCM-30;maddr=192.168.10.30
Codec-name: G722
Payload type: 9
```

Note If calls are not going through the Cisco Intercompany Media Engine, you can also use the **show tls-proxy session** command to troubleshoot the success of the TLS handshake between the components in the Cisco Intercompany Media Engine system. See the command reference for information about this command.

show uc-ime signaling-sessions statistics

Displays statistical information about corresponding signaling sessions stored by Cisco Intercompany Media Engine Proxy. Failure of signaling sessions in the Cisco Intercompany Media Engine can occur for different call-related reasons; such as failure of ticket verification or domain name verification, or offering RTP over the Internet.

```
hostname# show uc-ime signaling-sessions statistics
10 in use, 20 most used
15 terminated
  Ticket integrity check failed: 2
  Ticket decode failed: 1
  Ticket epoch mismatch: 1
  Ticket DID mismatch: 0
  Ticket timestamp invalid: 4
  Ticket domain check failed: 2
  Ticket not found: 0
  Route domain name check failed: 1
  RTP over UC-IME: 2
```

Note Call-related failures, for example, can be due to the service policy rule being reconfigured or the primary ASA operating in failover mode. If a service policy rule for the Cisco Intercompany Media Engine Proxy is removed (by using the **no service policy** command) and reconfigured, the first call trasversing the ASA will fail. To resolve this issue, you must additionally enter the **clear connection** command and restart the ASA. If the failure is due to failover, the connections from the primary ASA are not synchronized to the standby ASA.

show uc-ime media-sessions detail

Displays the details about all active media sessions (calls) stored for the Cisco Intercompany Media Engine Proxy. Use this command to display output from successful calls. Additionally, use this command to troubleshoot problems with IP phone audio, such as one-way audio. If no calls are currently up, this output will be blank.

```
hostname(config)# show uc-ime media-sessions detail
 2 in use, 5 most used
 Media-session: 10.194.108.109/21558 :: client ip 192.168.10.51/19520
 Call ID: ab6d7980-a7d11b08-50-1e0aa8c0@192.168.10.30
 Session ID: 2948-32325449-0@81a985c9-f3a1-55a0-3b19-96549a027259
   Lcl SRTP conn 10.194.108.109/21558 to 10.194.108.119/29824 tx_pkts 20203 rx_pkts 20200
   refcnt 3 : created by Inspect SIP, passthrough not set
   RTP monitoring is enabled
      Failover_state               :  0
      Sum_all_packets              :  20196
      Codec_payload_format         :  9
      RTP_ptime_ms                 :  20
      Max_RBLR_pct_x100            :  0
      Max_ITE_count_in_8_sec       :  0
```

```
      Max_BLS_ms                    :  0
      Max_PDV_usec                  :  1000
      Min_PDV_usec                  :  0
      Mov_avg_PDV_usec              :  109
      Total_ITE_count               :  0
      Total_sec_count               :  403
      Concealed_sec_count           :  0
      Severely_concealed_sec_count  :  0
      Max_call_interval_ms          :  118
      Total_SequenceNumber_Resets   :  0
 Media-session: 192.168.10.3/30930 :: client ip 10.194.108.119/29824
 Call ID: N/A
   Lcl RTP conn 192.168.10.3/30930 to 192.168.10.51/19520 tx_pkts 20201 rx_pkts 20203
```

show uc-ime fallback-notification statistics

Displays statistics about the PSTN fallback notifications to the Cisco UMC. Even if a call is sent over VoIP because the quality of the connection was good, the connection quality might worsen mid-call. To ensure an overall good experience for the end user, Cisco Intercompany Media Engine attempts to perform a mid-call fallback. Performing a mid-call fallback requires the adaptive security appliance to monitor the RTP packets coming from the Internet. If fallback is required, the adaptive security appliance sends a REFER message to Cisco UCM to tell it that it needs to fallback the call to PSTN.

Cisco Intercompany Media Engine uses a configurable hold-down timer to set the amount of time that adaptive security appliance waits before notifying Cisco UCM whether to fall back to PSTN.

```
hostname# show uc-ime fallback-notification statistics
UCM address: 172.23.32.37
 Total Notifications Sent: 10
```

show uc-ime mapping-service-sessions

When the Cisco Intercompany Media Engine Proxy is configured for an off path deployment, displays mapping-service requests and replies between the proxy and the local Cisco UMC. A TCP port on the ASA is configured to listen for mapping requests.

The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

```
Hostname# show uc-b2blink mapping-service-sessions
Total active sessions:  2
Session client (IP:Port)      Idle time
192.168.1.10:2001             0:01:01
192.168.1.20:3001             0:10:20
```

show uc-ime mapping-service-sessions statistics

Displays statistical information about the Cisco Intercompany Media Engine Proxy mapping service used in off path signaling.

```
Hostname# show uc-ime mapping-service-sessions statistics
Total active sessions:  2
Session client      Total       Responses   Failed      Pending     Idle
(IP:Port)           requests    sent        requests    responses   time
192.168.1.10:2001   10          9           1           0           0:01:01
192.168.1.20:3001   19          19          0           0           0:10:20
```

Feature History for Cisco Intercompany Media Engine Proxy

Table 69-1 lists the release history for this feature.

Table 69-1 Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Intercompany Media Engine Proxy	8.3(1)	The Cisco Intercompany Media Engine Proxy was introduced. The following pane was added to the ASDM: Configuration > Firewall > Unified Communications > UC-IME Proxy The following wizard was added to ASDM, which allows you to configure the Unified Communication proxies (including the Cisco Intercompany Media Engine Proxy): Wizards > Unified Communications Wizard

PART 12

Configuring Connection Settings and QoS

Configuring Connection Settings

This chapter describes how to configure connection settings for connections that go through the ASA, or for management connections, that go to the ASA. Connection settings include:

- Maximum connections (TCP and UDP connections, embryonic connections, per-client connections)
- Connection timeouts
- Dead connection detection
- TCP sequence randomization
- TCP normalization customization
- TCP state bypass
- Global timeouts

This chapter includes the following sections:

- Information About Connection Settings, page 71-1
- Licensing Requirements for Connection Settings, page 71-4
- Guidelines and Limitations, page 71-5
- Default Settings, page 71-5
- Configuring Connection Settings, page 71-5
- Feature History for Connection Settings, page 71-11

Information About Connection Settings

This section describes why you might want to limit connections and includes the following topics:

- TCP Intercept and Limiting Embryonic Connections, page 71-2
- Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility, page 71-2
- Dead Connection Detection (DCD), page 71-2
- TCP Sequence Randomization, page 71-3
- TCP Normalization, page 71-3
- TCP State Bypass, page 71-3

TCP Intercept and Limiting Embryonic Connections

Limiting the number of embryonic connections protects you from a DoS attack. The ASA uses the per-client limits and the embryonic connection limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the ASA acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the ASA receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

Note When you use TCP SYN cookie protection to protect servers from SYN attacks, you must set the embryonic connection limit lower than the TCP SYN backlog queue on the server that you want to protect. Otherwise, valid clients can nolonger access the server during a SYN attack.

To view TCP Intercept statistics, including the top 10 servers under attack, see Chapter 77, “Configuring Threat Detection.”

Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility

By default, TCP management connections have TCP Intercept always enabled. When TCP Intercept is enabled, it intercepts the 3-way TCP connection establishment handshake packets and thus deprives the ASA from processing the packets for clientless SSL. Clientless SSL requires the ability to process the 3-way handshake packets to provide selective ACK and other TCP options for clientless SSL connections. To disable TCP Intercept for management traffic, you can set the embryonic connection limit; only after the embryonic connection limit is reached is TCP Intercept enabled.

Dead Connection Detection (DCD)

DCD detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist.

When you enable DCD, idle timeout behavior changes. With idle timeout, DCD probes are sent to each of the two end-hosts to determine the validity of the connection. If an end-host fails to respond after probes are sent at the configured intervals, the connection is freed, and reset values, if configured, are sent to each of the end-hosts. If both end-hosts respond that the connection is valid, the activity timeout is updated to the current time and the idle timeout is rescheduled accordingly.

Enabling DCD changes the behavior of idle-timeout handling in the TCP normalizer. DCD probing resets the idle timeout on the connections seen in the **show conn** command. To determine when a connection that has exceeded the configured timeout value in the timeout command but is kept alive due to DCD probing, the **show service-policy** command includes counters to show the amount of activity from DCD.

TCP Sequence Randomization

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from 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 ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

TCP Normalization

The TCP normalization feature identifies abnormal packets that the ASA can act on when they are detected; for example, the ASA can allow, drop, or clear the packets. TCP normalization helps protect the ASA from attacks. TCP normalization is always enabled, but you can customize how some features behave.

The TCP normalizer includes non-configurable actions and configurable actions. Typically, non-configurable actions that drop or clear connections apply to packets that are always bad. Configurable actions (as detailed in "Customizing the TCP Normalizer with a TCP Map" section on page 71-6) might need to be customized depending on your network needs.

See the following guidelines for TCP normalization:

- The normalizer does not protect from SYN floods. The ASA includes SYN flood protection in other ways.
- The normalizer always sees the SYN packet as the first packet in a flow unless the ASA is in loose mode due to failover.

TCP State Bypass

By default, all traffic that goes through the ASA is inspected using the Adaptive Security Algorithm and is either allowed through or dropped based on the security policy. The ASA maximizes the firewall performance by checking the state of each packet (is this a new connection or an established

connection?) and assigning it to either the session management path (a new connection SYN packet), the fast path (an established connection), or the control plane path (advanced inspection). See the "Stateful Inspection Overview" section on page 1-24 for more detailed information about the stateful firewall.

TCP packets that match existing connections in the fast path can pass through the ASA without rechecking every aspect of the security policy. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), can stand in the way of asymmetrical routing solutions: both the outbound and inbound flow of a connection must pass through the same ASA.

For example, a new connection goes to ASA 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through ASA 1, then the packets will match the entry in the fast path, and are passed through. But if subsequent packets go to ASA 2, where there was not a SYN packet that went through the session management path, then there is no entry in the fast path for the connection, and the packets are dropped. Figure 71-1 shows an asymmetric routing example where the outbound traffic goes through a different ASA than the inbound traffic:

Figure 71-1 Asymmetric Routing

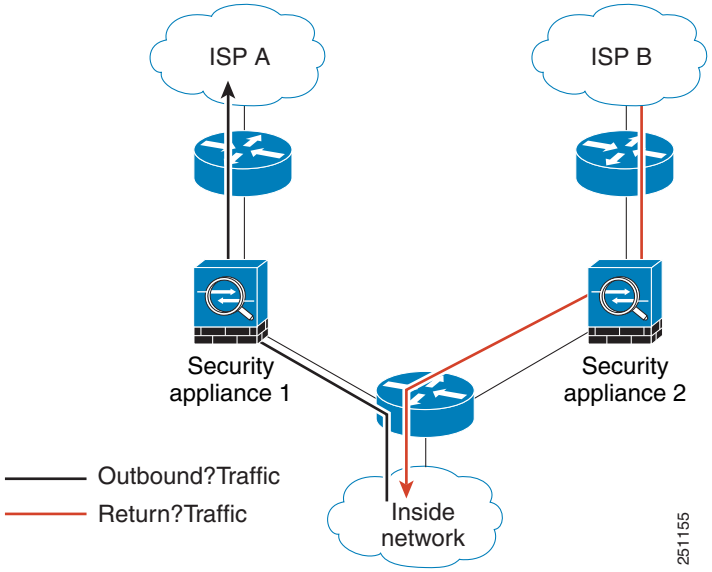

If you have asymmetric routing configured on upstream routers, and traffic alternates between two ASAs, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic much as it treats a UDP connection: when a non-SYN packet matching the specified networks enters the ASA, and there is not an fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Licensing Requirements for Connection Settings

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the following guidelines and limitations:

- TCP State Bypass Guidelines and Limitations, page 71-5

TCP State Bypass Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent mode.

Failover Guidelines

Failover is supported.

Unsupported Features

The following features are not supported when you use TCP state bypass:

- Application inspection—Application inspection requires both inbound and outbound traffic to go through the same ASA, so application inspection is not supported with TCP state bypass.
- AAA authenticated sessions—When a user authenticates with one ASA, traffic returning via the other ASA will be denied because the user did not authenticate with that ASA.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The ASA does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- SSM and SSC functionality—You cannot use TCP state bypass and any application running on an SSM or SSC, such as IPS or CSC.

NAT Guidelines

Because the translation session is established separately for each ASA, be sure to configure static NAT on both ASAs for TCP state bypass traffic; if you use dynamic NAT, the address chosen for the session on ASA 1 will differ from the address chosen for the session on ASA 2.

Default Settings

TCP State Bypass

TCP state bypass is disabled by default.

Configuring Connection Settings

This section includes the following topics:

- Customizing the TCP Normalizer with a TCP Map, page 71-6
- Configuring Connection Settings, page 71-8
- Configuring Global Timeouts, page 71-9

Task Flow For Configuring Configuration Settings (Except Global Timeouts)

Step 1 For TCP normalization customization, create a TCP map according to the "Customizing the TCP Normalizer with a TCP Map" section on page 71-6.

Step 2 For all connection settings except for global timeouts, configure a service policy according to Chapter 44, "Configuring a Service Policy."

Step 3 Configure connection settings according to the "Configuring Connection Settings" section on page 71-8.

Customizing the TCP Normalizer with a TCP Map

To customize the TCP normalizer, first define the settings using a TCP map.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Objects > TCP Maps** pane, and click **Add**.

The Add TCP Map dialog box appears.

Step 2 In the TCP Map Name field, enter a name.

Step 3 In the Queue Limit field, enter the maximum number of out-of-order packets, between 0 and 250 packets.

The Queue Limit sets the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection. The default is 0, which means this setting is disabled and the default system queue limit is used depending on the type of traffic:

- Connections for application inspection, IPS, and TCP check-retransmission have a queue limit of 3 packets. If the ASA receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.
- For other TCP connections, out-of-order packets are passed through untouched.

If you set the Queue Limit to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For example, for application inspection, IPS, and TCP check-retransmission traffic, any advertised settings from TCP packets are ignored in favor of the Queue Limit setting. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

Step 4 In the Timeout field, set the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds.

If they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the Queue Limit is set to 0; you need to set the limit to be 1 or above for the Timeout to take effect.

Step 5 In the Reserved Bits area, click **Clear and allow**, **Allow only**, or **Drop**.

Allow only allows packets with the reserved bits in the TCP header.

Clear and allow clears the reserved bits in the TCP header and allows the packet.

Drop drops the packet with the reserved bits in the TCP header.

Step 6 Check any of the following options:

- Clear urgent flag—Clears the URG flag through the ASA. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.
- Drop connection on window variation—Drops a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, "shrinking the window" is strongly discouraged. When this condition is detected, the connection can be dropped.
- Drop packets that exceed maximum segment size—Drops packets that exceed MSS set by peer.
- Check if transmitted data is the same as original—Enables the retransmit data checks.
- Drop packets which have past-window sequence—Drops packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window. If you do not check this option, then the Queue Limit must be set to 0 (disabled).
- Drop SYN Packets with data—Drops SYN packets with data.
- Enable TTL Evasion Protection—Enables the TTL evasion protection offered by the ASA. Do not enable this option if you want to prevent attacks that attempt to evade security policy.
- For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the ASA and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the ASA to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.
- Verify TCP Checksum—Enables checksum verification.
- Drop SYNACK Packets with data—Drops TCP SYNACK packets that contain data.
- Drop packets with invalid ACK—Drops packets with an invalid ACK. You might see invalid ACKs in the following instances:
  - In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.
  - Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.

Note TCP packets with an invalid ACK are automatically allowed for WAAS connections.

Step 7 To set TCP options, check any of the following options:

- Clear Selective Ack—Sets whether the selective-ack TCP option is allowed or cleared.
- Clear TCP Timestamp—Sets whether the TCP timestamp option is allowed or cleared.
- Clear Window Scale—Sets whether the window scale timestamp option is allowed or cleared.
- Range—Sets the valid TCP options ranges, which should fall within 6-7 and 9-255. The lower bound should be less than or equal to the upper bound. Choose **Allow** or **Drop** for each range.

Step 8 Click **OK**.

Configuring Connection Settings

To set connection settings, perform the following steps.

Guidelines and Limitations

Depending on the number of CPU cores on your ASA model, the maximum concurrent and embryonic connections may exceed the configured numbers due to the way each core manages connections. In the worst case scenario, the ASA allows up to *n*-1 extra connections and embryonic connections, where *n* is the number of cores. For example, if your model has 4 cores, if you configure 6 concurrent connections and 4 embryonic connections, you could have an additional 3 of each type. To determine the number of cores for your model, enter the **show cpu core** command.

Detailed Steps

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 44, "Configuring a Service Policy."

You can configure connection limits as part of a new service policy rule, or you can edit an existing service policy.

Step 2 On the Rule Actions dialog box, click the **Connection Settings** tab.

Step 3 To set maximum connections, configure the following values in the Maximum Connections area:

- TCP & UDP Connections—Specifies the maximum number of simultaneous TCP and UDP connections for all clients in the traffic class, up to 2000000. The default is 0 for both protocols, which means the maximum possible connections are allowed.
- Embryonic Connections—Specifies the maximum number of embryonic connections per host up to 2000000. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.
- Per Client Connections—Specifies the maximum number of simultaneous TCP and UDP connections for each client up to 2000000. When a new connection is attempted by a client that already has opened the maximum per-client number of connections, the ASA rejects the connection and drops the packet.
- Per Client Embryonic Connections—Specifies the maximum number of simultaneous TCP embryonic connections for each client up to 2000000. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the ASA, the ASA proxies the request to the TCP Intercept feature, which prevents the connection.

Step 4 To configure connection timeouts, configure the following values in the TCP Timeout area:

- Connection Timeout—Specifies the idle time until a connection slot (of *any* protocol, not just TCP) is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
- Send reset to TCP endpoints before timeout—Specifies that the ASA should send a TCP reset message to the endpoints of the connection before freeing the connection slot.
- Embryonic Connection Timeout—Specifies the idle time until an embryonic 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. See the "Customizing the TCP Normalizer with a TCP Map" section on page 71-6.

Step 7 Click **OK**.

Step 8 To set the time to live, check **Decrement time to live for a connection**.

Step 9 To enable TCP state bypass, in the Advanced Options area, check **TCP State Bypass**.

Step 10 Click **OK** or **Finish**.

Configuring Global Timeouts

The Configuration > Firewall > Advanced > Global Timeouts pane lets you set the timeout durations for use with the ASA. All durations are displayed in the format *hh*:*mm*:*ss*. It sets the idle time for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool. TCP connection slots are freed approximately 60 seconds after a normal connection close sequence.

Fields

In all cases, except for Authentication absolute and Authentication inactivity, unchecking the check boxes means there is no timeout value. For those two cases, clearing the check box means to reauthenticate on every new connection.

- Connection—Modifies the idle time until a connection slot is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.

- Half-closed—Modifies the idle time until a TCP half-closed connection closes. The minimum is 5 minutes. The default is 10 minutes. Enter 0:0:0 to disable timeout for a half-closed connection.
- UDP—Modifies the idle time until a UDP protocol connection closes. This duration must be at least 1 minute. The default is 2 minutes. Enter 0:0:0 to disable timeout.
- ICMP—Modifies the idle time after which general ICMP states are closed.
- H.323—Modifies the idle time until an H.323 media connection closes. The default is 5 minutes. Enter 0:0:0 to disable timeout.
- H.225—Modifies the idle time until an H.225 signaling connection closes. The H.225 default timeout is 1 hour (1:0:0). Setting the value of 0:0:0 means never close this connection. To close this connection immediately after all calls are cleared, a value of 1 second (0:0:1) is recommended.
- MGCP—Modifies the timeout value for MGCP which represents the idle time after which MGCP media ports are closed. The MGCP default timeout is 5 minutes (0:5:0). Enter 0:0:0 to disable timeout.
- MGCP PAT—Modifies the idle time after which an MGCP PAT translation is removed. The default is 5 minutes (0:5:0). The minimum time is 30 seconds. Uncheck the check box to return to the default value.
- TCP Proxy Reassembly—Configures the idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).
- Floating Connection—When multiple static routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To take advantage of this feature, change the timeout to a new value between 0:1:0 and 1193:0:0.
- SUNRPC—Modifies the idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes. Enter 0:0:0 to disable timeout.
- SIP—Modifies the idle time until an SIP signalling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- SIP Media—Modifies the idle time until an SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- SIP Provisional Media—Modifies the timeout value for SIP provisional media connections, between 0:1:0 and 1193:0:0. The default is 2 minutes.
- SIP Invite—Modifies the idle time after which pinholes for PROVISIONAL responses and media xlates will be closed. The minimum value is 0:1:0, the maximum value is 0:30:0. The default value is 0:3:0.
- SIP Disconnect—Modifies the idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message. The minimum value is 0:0:1, the maximum value is 0:10:0. The default value is 0:2:0.
- Authentication absolute—Modifies the duration until the authentication cache times out and you have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value. The system waits until you start a new connection to prompt you again. Enter 0:0:0 to disable caching and reauthenticate on every new connection.

Note Do not set this value to 0:0:0 if passive FTP is used on the connections.

Note When Authentication Absolute = 0, HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. This workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

- Authentication inactivity—Modifies the idle time until the authentication cache times out and users have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value.
- Translation Slot—Modifies the idle time until a translation slot is freed. This duration must be at least 1 minute. The default is 3 hours. Enter 0:0:0 to disable the timeout.
- (8.4(3) and later, not including 8.5(1) and 8.6(1)) PAT Translation Slot—Modifies the idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.

Feature History for Connection Settings

Table 71-1 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 71-1 Feature History for Connection Settings

Feature Name	Platform Releases	Feature Information
TCP state bypass	8.2(1)	This feature was introduced. The following command was introduced: **set connection advanced-options tcp-state-bypass**.
Connection timeout for all protocols	8.2(2)	The idle timeout was changed to apply to all protocols, not just TCP. The following screen was modified: Configuration > Firewall > Service Policies > Rule Actions > Connection Settings.
Timeout for connections using a backup static route	8.2(5)/8.4(2)	When multiple static routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To take advantage of this feature, change the timeout to a new value. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts.

Table 71-1 Feature History for Connection Settings (continued)

Feature Name	Platform Releases	Feature Information
Configurable timeout for PAT xlate	8.4(3)	When a PAT xlate times out (by default after 30 seconds), and the ASA reuses the port for a new translation, some upstream routers might reject the new connection because the previous connection might still be open on the upstream device. The PAT xlate timeout is now configurable, to a value between 30 seconds and 5 minutes. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts. *This feature is not available in 8.5(1) or 8.6(1).*
Increased maximum connection limits for service policy rules	9.0(1)	The maximum number of connections for service policy rules was increased from 65535 to 2000000. We modified the following screen: Configuration > Firewall > Service Policy Rules > Connection Settings.

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.

Note For the ASASM, we suggest performing QoS on the switch instead of the ASASM. Switches have more capability in this area.

This chapter describes how to apply QoS policies and includes the following sections:

- Information About QoS, page 72-1
- Licensing Requirements for QoS, page 72-5
- Guidelines and Limitations, page 72-5
- Configuring QoS, page 72-6
- Monitoring QoS, page 72-11
- Feature History for QoS, page 72-14

Information About QoS

You should consider that in an ever-changing network environment, QoS is not a one-time deployment, but an ongoing, essential part of network design.

This section describes the QoS features supported by the ASA and includes the following topics:

- Supported QoS Features, page 72-2
- What is a Token Bucket?, page 72-2
- Information About Policing, page 72-3
- Information About Priority Queuing, page 72-3
- Information About Traffic Shaping, page 72-4
- DSCP and DiffServ Preservation, page 72-5

Supported QoS Features

The ASA supports the following QoS features:

- Policing—To prevent individual flows from hogging the network bandwidth, you can limit the maximum bandwidth used per flow. See the "Information About Policing" section on page 72-3 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 "Information About Priority Queuing" section on page 72-3 for more information.
- Traffic shaping—If you have a device that transmits packets at a high speed, such as a ASA with Fast Ethernet, and it is connected to a low speed device such as a cable modem, then the cable modem is a bottleneck at which packets are frequently dropped. To manage networks with differing line speeds, you can configure the ASA to transmit packets at a fixed slower rate. See the "Information About Traffic Shaping" section on page 72-4 for more information.

What is a Token Bucket?

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer or a traffic shaper. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator.

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, an average rate, and a time interval. Although the average rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

average rate = burst size / time interval

Here are some definitions of these terms:

- Average rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.
- Burst size—Also called the Committed Burst (Bc) size, it specifies in bits or bytes per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For traffic shaping, it specifies bits per burst; for policing, it specifies bytes per burst.)
- Time interval—Also called the measurement interval, it specifies the time quantum in seconds per burst.

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of traffic shaping) or the packet is discarded or marked down (in the case of policing). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, or queue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the burstiness is bounded so that the flow will never send faster than the token bucket capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula:

(token bucket capacity in bits / time interval in seconds) + established rate in bps = maximum flow speed in bps

This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.

Information About Policing

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow or class can take over the entire resource. When traffic exceeds the maximum rate, the ASA drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Information About Priority Queuing

LLQ priority queuing lets you prioritize certain traffic flows (such as latency-sensitive traffic like voice and video) ahead of other traffic.

The ASA supports two types of priority queuing:

- Standard priority queuing—Standard priority queuing uses an LLQ priority queue on an interface (see the "Configuring the Standard Priority Queue for an Interface" section on page 72-8), while all other traffic goes into the "best effort" queue. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped. This is called *tail drop*. To avoid having the queue fill up, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of the priority queuing. Packets in the LLQ queue are always transmitted before packets in the best effort queue.
- Hierarchical priority queuing—Hierarchical priority queuing is used on interfaces on which you enable a traffic shaping queue. A subset of the shaped traffic can be prioritized. The standard priority queue is not used. See the following guidelines about hierarchical priority queuing:
  - Priority packets are always queued at the head of the shape queue so they are always transmitted ahead of other non-priority queued packets.
  - Priority packets are never dropped from the shape queue unless the sustained rate of priority traffic exceeds the shape rate.
  - For IPsec-encrypted packets, you can only match traffic based on the DSCP or precedence setting.
  - IPsec-over-TCP is not supported for priority traffic classification.

Information About Traffic Shaping

Traffic shaping is used to match device and link speeds, thereby controlling packet loss, variable delay, and link saturation, which can cause jitter and delay.

Note Traffic shaping is 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 or ASASM, 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 72-2.
- When bursty traffic exceeds the specified shape rate, packets are queued and transmitted later. Following are some characteristics regarding the shape queue (for information about hierarchical priority queuing, see the "Information About Priority Queuing" section on page 72-3):
  - The queue size is calculated based on the shape rate. The queue can hold the equivalent of 200-milliseconds worth of shape rate traffic, assuming a 1500-byte packet. The minimum queue size is 64.
  - When the queue limit is reached, packets are tail-dropped.
  - Certain critical keep-alive packets such as OSPF Hello packets are never dropped.
  - The time interval is derived by *time_interval* = *burst_size* / *average_rate*. The larger the time interval is, the burstier the shaped traffic might be, and the longer the link might be idle. The effect can be best understood using the following exaggerated example:

    Average Rate = 1000000

    Burst Size = 1000000

    In the above example, the time interval is 1 second, which means, 1 Mbps of traffic can be bursted out within the first 10 milliseconds of the 1-second interval on a 100 Mbps FE link and leave the remaining 990 milliseconds idle without being able to send any packets until the next time interval. So if there is delay-sensitive traffic such as voice traffic, the Burst Size should be reduced compared to the average rate so the time interval is reduced.

How QoS Features Interact

You can configure each of the QoS features alone if desired for the ASA. Often, though, you configure multiple QoS features on the ASA so you can prioritize some traffic, for example, and prevent other traffic from causing bandwidth problems.

See the following supported feature combinations per interface:

- Standard priority queuing (for specific traffic) + Policing (for the rest of the traffic).

  You cannot configure priority queuing and policing for the same set of traffic.
- Traffic shaping (for all traffic on an interface) + Hierarchical priority queuing (for a subset of traffic).

You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed. For example, if you configure standard priority queuing for the global policy, and then configure traffic shaping for a specific interface, the feature you configured last is rejected because the global policy overlaps the interface policy.

Typically, if you enable traffic shaping, you do not also enable policing for the same traffic, although the ASA does not restrict you from configuring this.

DSCP and DiffServ Preservation

- DSCP markings are preserved on all traffic passing through the ASA.
- The ASA does not locally mark/remark any classified traffic, but it honors the Expedited Forwarding (EF) DSCP bits of every packet to determine if it requires “priority” handling and will direct those packets to the LLQ.
- DiffServ marking is preserved on packets when they traverse the service provider backbone so that QoS can be applied in transit (QoS tunnel pre-classification).

Licensing Requirements for QoS

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single context mode only. Does not support multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode only. Does not support transparent firewall mode.

IPv6 Guidelines

Does not support IPv6.

Model Guidelines

- (Multi-processor models, such as the ASA 5500-X, ASA 5580, ASA 5585-X, and ASASM) Traffic shaping is not supported.
- (ASA 5580) You cannot create a standard priority queue for a Ten Gigabit Ethernet interface. **Note**: For the ASA 5585-X, standard priority queuing is supported on a Ten Gigabit Interface.

- (ASA 5512-X through ASA 5555-X) Priority queuing is not supported on the Management 0/0 interface.
- (ASASM) Only policing is supported.

Additional Guidelines and Limitations

- QoS is applied unidirectionally; only traffic that enters (or exits, depending on the QoS feature) the interface to which you apply the policy map is affected. See the "Feature Directionality" section on page 38-3 for more information.
- For traffic shaping, you can only use the **class-default** class map, which is automatically created by the ASA, and which matches all traffic.
- For priority traffic, you cannot use the **class-default** class map.
- For hierarchical priority queuing, for encrypted VPN traffic, you can only match traffic based on the DSCP or precedence setting; you cannot match a tunnel group.
- For hierarchical priority queuing, IPsec-over-TCP traffic is not supported.
- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed.
- For standard priority queuing, the queue must be configured for a physical interface or, for the ASA 5505 or ASASM, a VLAN.
- For policing, to-the-box traffic is not supported.
- For policing, traffic to and from a VPN tunnel bypass interface is not supported.
- For policing, when you match a tunnel group class map, only outbound policing is supported.

Configuring QoS

This section includes the following topics:

- Determining the Queue and TX Ring Limits for a Standard Priority Queue, page 72-7
- Configuring the Standard Priority Queue for an Interface, page 72-8
- Configuring a Service Rule for Standard Priority Queuing and Policing, page 72-9
- Configuring a Service Rule for Traffic Shaping and Hierarchical Priority Queuing, page 72-10

Determining the Queue and TX Ring Limits for a Standard Priority Queue

To determine the priority queue and TX ring limits, use the worksheets below.

Table 72-1 shows how to calculate the priority queue size. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can adjust the queue buffer size according to the "Configuring the Standard Priority Queue for an Interface" section on page 72-8.

Table 72-1 Queue Limit Worksheet

Step 1	_____ *Outbound bandwidth (Mbps or Kbps)*[1]	Mbps	x	**125**	=	_____ *# of bytes/ms*		
		Kbps	x	**.125**	=	_____ *# of bytes/ms*		
Step 2	_____ *# of bytes/ms from Step 1*	÷	_____ *Average packet size (bytes)*[2]	x	_____ *Delay (ms)*[3]	=	_____ *Queue limit (# of packets)*	

1. For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.
2. Determine this value from a codec or sampling size. For example, for VoIP over VPN, you might use 160 bytes. We recommend 256 bytes if you do not know what size to use.
3. The delay depends on your application. For example, the recommended maximum delay for VoIP is 200 ms. We recommend 500 ms if you do not know what delay to use.

Table 72-2 shows how to calculate the TX ring limit. This limit determines the maximum number of packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears. This setting guarantees that the hardware-based transmit ring imposes a limited amount of extra latency for a high-priority packet.

Table 72-2 TX Ring Limit Worksheet

Step 1	_____ *Outbound bandwidth (Mbps or Kbps)*[1]	Mbps	x	**125**	=	_____ *# of bytes/ms*		
		Kbps	x	**0.125**	=	_____ *# of bytes/ms*		
Step 2	_____ *# of bytes/ms from Step 1*	÷	_____ *Maximum packet size (bytes)*[2]	x	_____ *Delay (ms)*[3]	=	_____ *TX ring limit (# of packets)*	

1. For example, DSL might have an uplink speed of 768 Kbps.Check with your provider.

2. Typically, the maximum size is 1538 bytes, or 1542 bytes for tagged Ethernet. If you allow jumbo frames (if supported for your platform), then the packet size might be larger.
3. The delay depends on your application. For example, to control jitter for VoIP, you should use 20 ms.

Configuring the Standard Priority Queue for an Interface

If you enable standard priority queuing for traffic on a physical interface, then you need to also create the priority queue on each interface. Each physical interface uses two queues: one for priority traffic, and the other for all other traffic. For the other traffic, you can optionally configure policing.

Note The standard priority queue is not required for hierarchical priority queuing with traffic shaping; see the "Information About Priority Queuing" section on page 72-3 for more information.

Restrictions

- (ASASM) The ASASM does not support priority queuing.
- (ASA 5580) You cannot create a standard priority queue for a Ten Gigabit Ethernet interface. **Note**: For the ASA 5585-X, standard priority queuing is supported on a Ten Gigabit Interface.
- (ASA 5512-X through ASA 5555-X) Priority queuing is not supported on the Management 0/0 interface.

Detailed Steps

Step 1 Go to Configuration > Device Management > Advanced > Priority Queue, and click **Add**.

The Add Priority Queue dialog box displays.

Step 2 From the Interface drop-down list, choose the physical interface name on which you want to enable the priority queue, or for the ASA 5505 or ASASM, the VLAN interface name.

Step 3 To change the size of the priority queues, in the Queue Limit field, enter the number of average, 256-byte packets that the specified interface can transmit in a 500-ms interval.

A packet that stays more than 500 ms in a network node might trigger a timeout in the end-to-end application. Such a packet can be discarded in each network node.

Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can use this option to increase the queue buffer size.

The upper limit of the range of values for this option is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Queue Limit that you specify affects both the higher priority low-latency queue and the best effort queue.

Step 4 To specify the depth of the priority queues, in the Transmission Ring Limit field, enter the number of maximum 1550-byte packets that the specified interface can transmit in a 10-ms interval.

This setting guarantees that the hardware-based transmit ring imposes no more than 10-ms of extra latency for a high-priority packet.

This option sets the maximum number of low-latency or normal priority packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears.

The upper limit of the range of values is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Transmission Ring Limit that you specify affects both the higher priority low-latency queue and the best-effort queue.

Configuring a Service Rule for Standard Priority Queuing and Policing

You can configure standard priority queuing and policing for different class maps within the same policy map. See the "How QoS Features Interact" section on page 72-4 for information about valid QoS configurations.

To create a policy map, perform the following steps.

Restrictions

- You cannot use the **class-default** class map for priority traffic.
- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed.
- (ASASM) The ASASM only supports policing.
- For policing, to-the-box traffic is not supported.
- For policing, traffic to and from a VPN tunnel bypass interface is not supported.
- For policing, when you match a tunnel group class map, only outbound policing is supported.

Guidelines

- For priority traffic, identify only latency-sensitive traffic.
- For policing traffic, you can choose to police all other traffic, or you can limit the traffic to certain types.

Detailed Steps

Step 1 To configure priority queuing, configure a service policy rule in the Configuration > Firewall > Service Policy Rules pane according to Chapter 44, "Configuring a Service Policy."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

Step 2 In the Rule Actions dialog box, click the **QoS** tab.

Step 3 Click **Enable priority for this flow**.

If this service policy rule is for an individual interface, ASDM automatically creates the priority queue for the interface (Configuration > Device Management > Advanced > Priority Queue; for more information, see the "Configuring the Standard Priority Queue for an Interface" section on page 72-8). If this rule is for the global policy, then you need to manually add the priority queue to one or more interfaces *before* you configure the service policy rule.

Step 4 Click **Finish**. The service policy rule is added to the rule table.

Step 5 To configure policing, configure a service policy rule for the same interface in the Configuration > Firewall > Service Policy Rules pane according to Chapter 44, "Configuring a Service Policy."

For policing traffic, you can choose to police all traffic that you are not prioritizing, or you can limit the traffic to certain types.

Step 6 In the Rule Actions dialog box, click the **QoS** tab.

Step 7 Click **Enable policing**, then check the **Input policing** or **Output policing** (or both) check boxes to enable the specified type of traffic policing. For each type of traffic policing, configure the following fields:

- Committed Rate—The rate limit for this traffic flow; this is a value in the range 8000-2000000000, specifying the maximum speed (bits per second) allowed.
- Conform Action—The action to take when the rate is less than the conform-burst value. Values are transmit or drop.
- Exceed Action—Take this action when the rate is between the conform-rate value and the conform-burst value. Values are transmit or drop.
- Burst Rate—A value in the range 1000-512000000, specifying the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value.

Step 8 Click **Finish**. The service policy rule is added to the rule table.

Step 9 Click **Apply** to send the configuration to the device.

Configuring a Service Rule for Traffic Shaping and Hierarchical Priority Queuing

You can configure traffic shaping for all traffic on an interface, and optionally hierarchical priority queuing for a subset of latency-sensitive traffic.

Guidelines

- One side-effect of priority queuing is packet re-ordering. For IPsec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings are false alarms in the case of priority queuing. You can configure the IPsec anti-replay window size to avoid possible false alarms. See the Configuration > VPN > IPsec > IPsec Rules > Enable Anti-replay window size option in the "Adding Crypto Maps" section on page 86-10.
- For hierarchical priority queuing, you do not need to create a priority queue on an interface.

Restrictions

- For hierarchical priority queuing, for encrypted VPN traffic, you can only match traffic based on the DSCP or precedence setting; you cannot match a tunnel group.
- For hierarchical priority queuing, IPsec-over-TCP traffic is not supported.
- Traffic shaping is not supported on the ASA 5580.
- For traffic shaping, you can only use the **class-default** class map, which is automatically created by the ASA, and which matches all traffic.

- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed. See the "How QoS Features Interact" section on page 72-4 for information about valid QoS configurations.
- You cannot configure traffic shaping in the global policy.

Detailed Steps

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 44, "Configuring a Service Policy."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

Step 2 In the Rule Actions dialog box, click the **QoS** tab.

Step 3 Click **Enable traffic shaping**, and configure the following fields:

- Average Rate—Sets the average rate of traffic in bits per second over a given fixed time period, between 64000 and 154400000. Specify a value that is a multiple of 8000.
- Burst Size—Sets the average burst size in bits that can be transmitted over a given fixed time period, between 2048 and 154400000. Specify a value that is a multiple of 128. If you do not specify the Burst Size, the default value is equivalent to 4-milliseconds of traffic at the specified Average Rate. For example, if the average rate is 1000000 bits per second, 4 ms worth = 1000000 * 4/1000 = 4000.

Step 4 (Optional) To configure priority queuing for a subset of shaped traffic:

a. Click **Enforce priority to selected shape traffic**.

b. Click **Configure** to identify the traffic that you want to prioritize.

You are prompted to identify the traffic for which you want to apply priority queuing.

c. After you identify the traffic (see the "Adding a Service Policy Rule for Through Traffic" section on page 44-8), click **Next**.

d. Click **Enable priority for this flow**.

e. Click **Finish**.

You return to the QoS tab.

Step 5 Click **Finish**. The service policy rule is added to the rule table.

Step 6 Click **Apply** to send the configuration to the device.

Monitoring QoS

To monitor QoS in ASDM, you can enter commands at the Command Line Interface tool. This section includes the following topics:

- Viewing QoS Police Statistics, page 72-12
- Viewing QoS Standard Priority Statistics, page 72-12
- Viewing QoS Shaping Statistics, page 72-12
- Viewing QoS Standard Priority Queue Statistics, page 72-13

Viewing QoS Police Statistics

To view the QoS statistics for traffic policing, use the **show service-policy** command with the **police** keyword:

```
hostname# show service-policy police
```

The following is sample output for the **show service-policy police** command:

```
hostname# show service-policy police

Global policy:
    Service-policy: global_fw_policy

Interface outside:
    Service-policy: qos
        Class-map: browse
            police Interface outside:
                cir 56000 bps, bc 10500 bytes
                conformed 10065 packets, 12621510 bytes; actions: transmit
                exceeded 499 packets, 625146 bytes; actions: drop
                conformed 5600 bps, exceed 5016 bps
        Class-map: cmap2
            police Interface outside:
                cir 200000 bps, bc 37500 bytes
                conformed 17179 packets, 20614800 bytes; actions: transmit
                exceeded 617 packets, 770718 bytes; actions: drop
                conformed 198785 bps, exceed 2303 bps
```

Viewing QoS Standard Priority Statistics

To view statistics for service policies implementing the **priority** command, use the **show service-policy** command with the **priority** keyword:

```
hostname# show service-policy priority
```

The following is sample output for the **show service-policy priority** command:

```
hostname# show service-policy priority
Global policy:
    Service-policy: global_fw_policy
Interface outside:
    Service-policy: qos
        Class-map: TG1-voice
            Priority:
                Interface outside: aggregate drop 0, aggregate transmit 9383
```

Note "Aggregate drop" denotes the aggregated drop in this interface; "aggregate transmit" denotes the aggregated number of transmitted packets in this interface.

Viewing QoS Shaping Statistics

To view statistics for service policies implementing the **shape** command, use the **show service-policy** command with the **shape** keyword:

```
hostname# show service-policy shape
```

The following is sample output for the **show service-policy shape** command:

```
hostname# show service-policy shape
Interface outside
  Service-policy: shape
    Class-map: class-default

      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

      shape (average) cir 2000000, bc 8000, be 8000
```

The following is sample output of the **show service policy shape** command, which includes service policies that include the **shape** command and the **service-policy** command that calls the hierarchical priority policy and the related statistics:

```
hostname# show service-policy shape

Interface outside:
  Service-policy: shape
    Class-map: class-default

      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

      shape (average) cir 2000000, bc 16000, be 16000

      Service-policy: voip
        Class-map: voip

          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
        Class-map: class-default

          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
```

Viewing QoS Standard Priority Queue Statistics

To display the priority-queue statistics for an interface, use the **show priority-queue statistics** command in privileged EXEC mode. The results show the statistics for both the best-effort (BE) queue and the low-latency queue (LLQ). The following example shows the use of the **show priority-queue statistics** command for the interface named test, and the command output.

```
hostname# show priority-queue statistics test

Priority-Queue Statistics interface test

Queue Type         = BE
Packets Dropped    = 0
Packets Transmit   = 0
Packets Enqueued   = 0
```

```
Current Q Length  = 0
Max Q Length      = 0

Queue Type        = LLQ
Packets Dropped   = 0
Packets Transmit  = 0
Packets Enqueued  = 0
Current Q Length  = 0
Max Q Length      = 0
hostname#
```

In this statistical report, the meaning of the line items is as follows:

- "Packets Dropped" denotes the overall number of packets that have been dropped in this queue.
- "Packets Transmit" denotes the overall number of packets that have been transmitted in this queue.
- "Packets Enqueued" denotes the overall number of packets that have been queued in this queue.
- "Current Q Length" denotes the current depth of this queue.
- "Max Q Length" denotes the maximum depth that ever occurred in this queue.

Feature History for QoS

Table 72-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 72-3 Feature History for QoS

Feature Name	Platform Releases	Feature Information
Priority queuing and policing	7.0(1)	We introduced QoS priority queuing and policing. We introduced the following screens: Configuration > Device Management > Advanced > Priority Queue Configuration > Firewall > Service Policy Rules
Shaping and hierarchical priority queuing	7.2(4)/8.0(4)	We introduced QoS shaping and hierarchical priority queuing. We modified the following screen: Configuration > Firewall > Service Policy Rules.
Ten Gigabit Ethernet support for a standard priority queue on the ASA 5585-X	8.2(3)/8.4(1)	We added support for a standard priority queue on Ten Gigabit Ethernet interfaces for the ASA 5585-X.

Troubleshooting Connections and Resources

This chapter describes how to troubleshoot the ASA and includes the following sections:

- Testing Your Configuration, page 73-1
- Monitoring Performance, page 73-8
- Monitoring System Resources, page 73-9
- Monitoring Connections, page 73-11
- Monitoring Per-Process CPU Usage, page 73-12

Testing Your Configuration

This section describes how to test connectivity for the single mode ASA or for each security context, how to ping the ASA interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

This section includes the following topics:

- Pinging ASA Interfaces, page 73-1
- Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping, page 73-3
- Determining Packet Routing with Traceroute, page 73-6
- Tracing Packets with Packet Tracer, page 73-7

Pinging ASA Interfaces

To test whether the ASA interfaces are up and running and that the ASA and connected routers are operating correctly, you can ping the ASA interfaces.

To ping the ASA interfaces, perform the following steps:

Step 1 Draw a diagram of your single-mode ASA or security context that shows the interface names, security levels, and IP addresses.

Note Although this procedure uses IP addresses, the **ping** command also supports DNS names and names that are assigned to a local IP address with the **name** command.

The diagram should also include any directly connected routers and a host on the other side of the router from which you will ping the ASA. (See Figure 73-1.)

Figure 73-1 Network Diagram with Interfaces, Routers, and Hosts

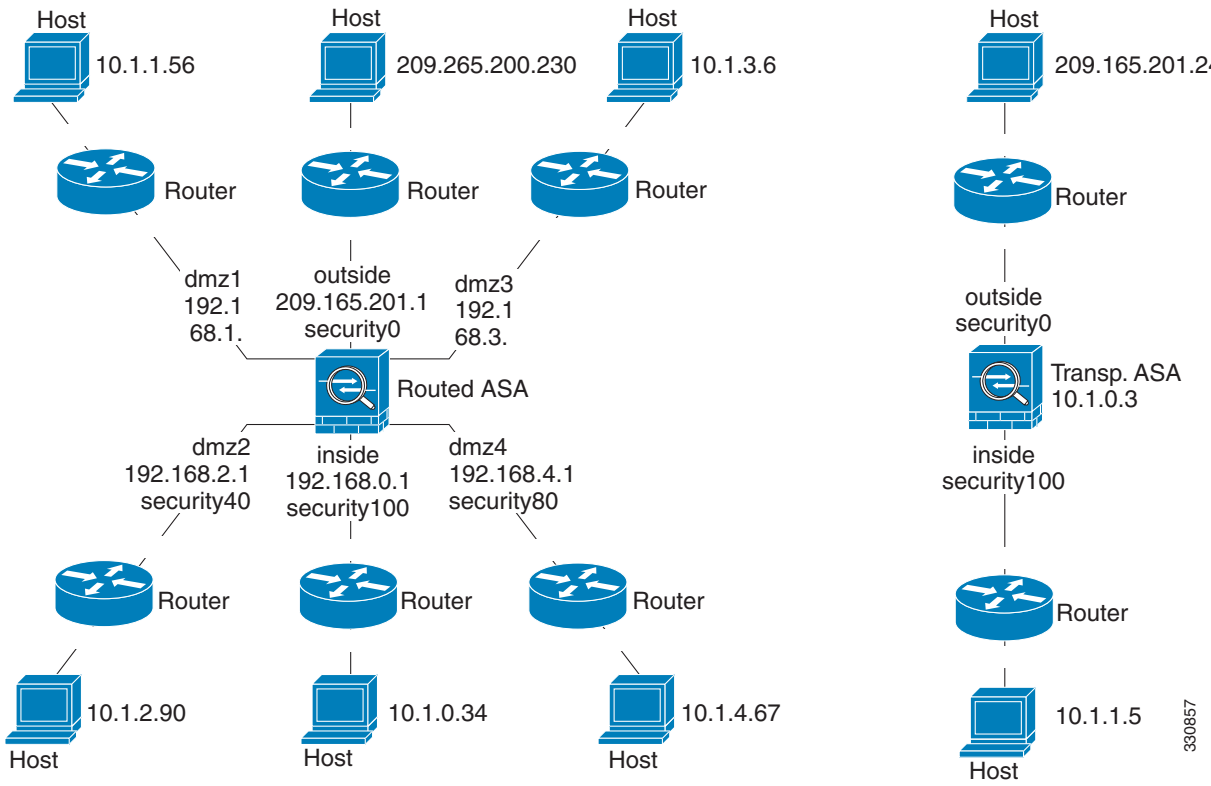

Step 2 Ping each ASA interface from the directly connected routers. For transparent mode, ping the management IP address. This test ensures that the ASA interfaces are active and that the interface configuration is correct.

A ping might fail if the ASA interface is not active, the interface configuration is incorrect, or if a switch between the ASA and a router is down (see Figure 73-2). In this case, no debugging messages or syslog messages appear, because the packet never reaches the ASA.

Figure 73-2 Ping Failure at the ASA Interface

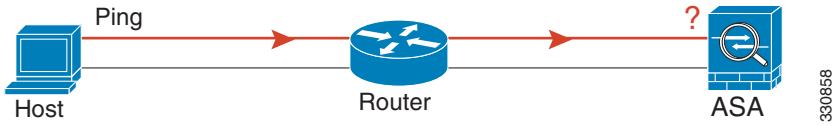

If the ping reaches the ASA, and it responds, debugging messages similar to the following appear:

```
ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
```

If the ping reply does not return to the router, then a switch loop or redundant IP addresses may exist (see Figure 73-3).

Figure 73-3 Ping Failure Because of IP Addressing Problems

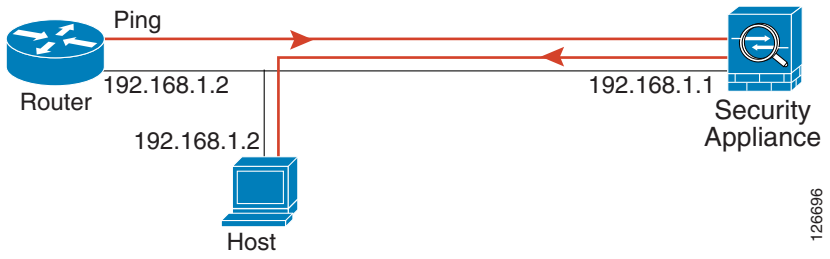

Step 3 Ping each ASA interface from a remote host. For transparent mode, ping the management IP address. This test checks whether the directly connected router can route the packet between the host and the ASA, and whether the ASA can correctly route the packet back to the host.

A ping might fail if the ASA does not have a return route to the host through the intermediate router (see Figure 73-4). In this case, the debugging messages show that the ping was successful, but syslog message 110001 appears, indicating a routing failure has occurred.

Figure 73-4 Ping Failure Because the ASA Has No Return Route

Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping

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

This section includes the following topics:

- Information About Ping, page 73-3
- Pinging From an ASA Interface, page 73-4
- Pinging to an ASA Interface, page 73-4
- Pinging Through the ASA Interface, page 73-4
- Troubleshooting the Ping Tool, page 73-4
- Using the Ping Tool, page 73-5

Information About Ping

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

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

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

- Loopback testing of two interfaces—A ping may be initiated from one interface to another on the same ASA, as an external loopback test to verify basic "up" status and operation of each interface.
- Pinging to an ASA—The Ping tool can ping an interface on another ASA to verify that it is up and responding.
- Pinging through an ASA—Ping packets originating from the Ping tool may pass through an intermediate ASA on their way to a device. The echo packets will also pass through two of its interfaces as they return. This procedure can be used to perform a basic test of the interfaces, operation, and response time of the intermediate unit.
- Pinging to test questionable operation of a network device—A ping may be initiated from an ASA interface to a network device that is suspected of functioning incorrectly. If the interface is configured correctly and an echo is not received, there may be problems with the device.
- Pinging to test intermediate communications—A ping may be initiated from an ASA interface to a network device that is known to be functioning correctly and returning echo requests. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

Pinging From an ASA Interface

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

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

Pinging to an ASA Interface

When you try to ping to an ASA interface, verify that the pinging response (ICMP echo reply*)* is enabled for that interface by choosing **Tools > Ping**. When pinging is disabled, the ASA cannot be detected by other devices or software applications, and does not respond to the ASDM Ping tool.

Pinging Through the ASA Interface

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

To enable internal hosts to ping external hosts, configure ICMP inspection. Choose **Configuration > Firewall > Service Policies**.

Troubleshooting the Ping Tool

When pings fail to receive an echo, it may be the result of a configuration or operational error in an ASA, and not necessarily because of no response from the IP address being pinged. Before using the Ping tool to ping from, to, or through an ASA interface, perform the following basic checks:

- Verify that interfaces are configured. Choose **Configuration > Device Setup > Interfaces**.

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

Using the Ping Tool

To use the Ping tool, perform the following steps:

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

The Ping dialog box appears.

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

Ping also supports IPv6 addresses.

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

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

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

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

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

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

If the ping fails, the output is as follows:

```
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
?????
Success rate is 0 percent (0/5)
```

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

Determining Packet Routing with Traceroute

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

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

To use the Traceroute tool, perform the following steps:

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

The Traceroute dialog box appears.

Step 2 Enter hostname or IP address to which the route is traced. If the hostname is given, define it by choosing **Configuration > Firewall > Objects > Service Objects/Groups**, or configure a DNS server to enable this tool to resolve the hostname to an IP address.

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

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

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

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

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

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

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

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

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

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

Tracing Packets with Packet Tracer

The packet tracer tool provides packet tracing for packet sniffing and network fault isolation, as well as detailed information about the packets and how they are processed by the ASA. If a configuration command did not cause the packet to drop, the packet tracer tool can provide information about the cause in an easily readable format.

In addition, you can trace the lifespan of a packet through the ASA to see whether the packet is operating correctly with the packet tracer tool. This tool enables you to do the following:

- Debug all packet drops in a production network.
- Verify the configuration is working as intended.
- Show all rules applicable to a packet, along with the CLI commands that caused the rule addition.
- Show a time line of packet changes in a data path.
- Inject tracer packets into the data path.
- Search for an IPv4 or IPv6 address based on the user identity and the FQDN.

To use the packet tracer, perform the following steps:

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

The Cisco ASDM Packet Tracer dialog box appears.

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

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

Step 4 In the Source drop-down list, select one of the following options:

- IP Address
- User
- FQDN
- Security Tag
- Security Name

  Select the Security Tag or Security Name options when you want to trace packets sent by the ASA when integrated with the Cisco TrustSec solution. Security names are created on the Cisco ISE and provide user-friendly names for security groups.

  If a security policy is configured on the ASA with that security tags or security names, the ASA enforces the policy. (You can create security policies on the ASA that contain security tags or security names. To enforce policies based on security group names, the ASA needs the security group table to map security names to security tags.)

  See Configuring the ASA to Integrate with Cisco TrustSec, page 48-1 for information about configuring the ASA to integrate with the Cisco TrustSec solution.

Step 5 Based on the option you selected from the Source drop-down list, enter the corresponding text for the item you want to trace; for example, enter the source IP address for the packet trace in the Source IP Address field.

Step 6 For TCP and UDP only, choose the source port for the packet trace from the drop-down list.

Step 7 In the Destination drop-down list, select one of the following options:

- IP Address
- FQDN

- Security Tag
- Security Name

Step 8 Based on the option you selected from the Destination drop-down list, enter the corresponding text for the item you want to trace; for example, enter the source IP address for the packet trace in the Destination IP Address field.

Step 9 For TCP and UDP only, choose the destination port for the packet trace from the drop-down list.

Step 10 For ICMP only, choose the type of packet trace from the Type drop-down list. Then enter the trace code and trace ID in the appropriate fields.

Step 11 For IP only, enter the protocol number in the Protocol field. Valid values range from 0 to 255.

Step 12 Click **Start** to trace the packet.

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

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

Step 13 Click **Clear** to start a new packet trace.

Monitoring Performance

To view ASA performance information in a graphical or tabular format, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > Connection Graphs > Perfmon**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- AAA Perfmon—Displays the ASA AAA performance information.
- Inspection Perfmon—Displays the ASA inspection performance information.
- Web Perfmon—Displays the ASA web performance information, including URL access and URL server requests.
- Connections Perfmon—Displays the ASA connections performance information.
- Xlate Perfmon—Displays the ASA NAT performance information.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view performance statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected performance statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the performance statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected performance statistics.

Monitoring System Resources

This section includes the following topics:

- Blocks, page 73-9
- CPU, page 73-10
- Memory, page 73-10

Blocks

To view the free and used memory blocks, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > Blocks**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- Blocks Used—Displays the ASA used memory blocks.
- Blocks Free—Displays the ASA free memory blocks.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected memory block statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the memory block statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected memory block statistics.

CPU

To view the CPU utilization, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > CPU**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected CPU utilization statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the CPU utilization statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected CPU utilization statistics.

Memory

To view the memory utilization, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > Blocks**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- Free Memory—Displays the ASA free memory.
- Used Memory—Displays the ASA used memory.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected memory utilization statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the memory utilization statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected memory utilization statistics.

Monitoring Connections

To view current connections in a tabular format, in the ASDM main window, choose **Monitoring > Properties > Connections**. Each connection is identified by the following parameters:

- Protocol
- Source:
  - Security ID
  - Security Name
  - IP address
  - Port
- Destination:
  - Security ID
  - Security Name
  - IP address
  - Port

- Idle time since the last packet was sent or received
- Amount of sent and received traffic on the connection

Monitoring Per-Process CPU Usage

You can monitor the processes that run on the CPU. You can obtain information about the percentage of CPU that is used by a certain process. CPU usage statistics are sorted in descending order to display the highest consumer at the top. Also included is information about the load on the CPU per process, at 5 seconds, 1 minute, and 5 minutes before the log time. This information is updated automatically every 5 seconds to provide real-time statistics. In ASDM, it is updated every 30 seconds.

To view CPU usage on a per-process basis, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > Per-Process CPU Usage**.

Step 2 To pause the auto-refresh of the screen, click **Stop auto-refresh**.

Step 3 To save the information on the screen to a local text file, click **Save log to local file**.

The Save dialog box appears.

Step 4 Enter the name of the text file, then click **Save**.

To color code processes according to their CPU usage range, click **Configure CPU usage**.

The Color Settings dialog box appears.

Step 5 Choose one of the following range options: 49% and below, 50% to 79%, and 80% and above.

Step 6 Click the foreground or background cell to display the Pick a Color dialog box, and select the foreground and background colors for the given ranges.

Step 7 Click one of the following tabs to pick the color palette: **Swatches**, **HSB**, or **RGB**. When you are done, click **OK**.

Step 8 Click **OK** to view the color-coded entries.

Step 9 Click **Refresh** to refresh the data manually at any time.

PART 13

Configuring Advanced Network Protection

CHAPTER 75

Configuring the ASA for Cisco Cloud Web Security

Cisco Cloud Web Security provides web security and web filtering services through the Software-as-a-Service (SaaS) model. Enterprises with the ASA in their network can use Cloud Web Security services without having to install additional hardware.

When Cloud Web Security is enabled on the ASA, the ASA transparently redirects selected HTTP and HTTPS traffic to the Cloud Web Security proxy servers. The Cloud Web Security proxy servers then scan the content and allow, block, or send a warning about the traffic based on the policy configured in Cisco ScanCenter to enforce acceptable use and to protect users from malware.

The ASA can optionally authenticate and identify users with Identity Firewall (IDFW) and AAA rules. The ASA encrypts and includes the user credentials (including usernames and/or user groups) in the traffic it redirects to Cloud Web Security. The Cloud Web Security service then uses the user credentials to match the traffic to the policy. It also uses these credentials for user-based reporting. Without user authentication, the ASA can supply an (optional) default username and/or group, although usernames and groups are not required for the Cloud Web Security service to apply policy.

You can customize the traffic you want to send to Cloud Web Security when you create your service policy rules. You can also configure a "whitelist" so that a subset of web traffic that matches the service policy rule instead goes directly to the originally requested web server and is not scanned by Cloud Web Security.

You can configure a primary and a backup Cloud Web Security proxy server, each of which the ASA polls regularly to check for availability.

Note This feature is also called "ScanSafe," so the ScanSafe name appears in some commands.

This chapter includes the following sections:

- Information About Cisco Cloud Web Security, page 75-2
- Licensing Requirements for Cisco Cloud Web Security, page 75-7
- Prerequisites for Cloud Web Security, page 75-7
- Guidelines and Limitations, page 75-7
- Default Settings, page 75-8
- Configuring Cisco Cloud Web Security, page 75-8
- Monitoring Cloud Web Security, page 75-26
- Related Documents, page 75-27
- Feature History for Cisco Cloud Web Security, page 75-27

Information About Cisco Cloud Web Security

This section includes the following topics:

- Redirection of Web Traffic to Cloud Web Security, page 75-2
- User Authentication and Cloud Web Security, page 75-2
- Authentication Keys, page 75-3
- ScanCenter Policy, page 75-4
- Cloud Web Security Actions, page 75-5
- Bypassing Scanning with Whitelists, page 75-5
- IPv4 and IPv6 Support, page 75-6
- Failover from Primary to Backup Proxy Server, page 75-6

Redirection of Web Traffic to Cloud Web Security

When an end user sends an HTTP or HTTPS request, the ASA receives it and optionally retrieves the user and/or group information. If the traffic matches an ASA service policy rule for Cloud Web Security, then the ASA redirects the request to the Cloud Web Security proxy servers. The ASA acts as an intermediary between the end user and the Cloud Web Security proxy server by redirecting the connection to the proxy server. The ASA changes the destination IP address and port in the client requests and adds Cloud Web Security-specific HTTP headers and then sends the modified request to the Cloud Web Security proxy server. The Cloud Web Security HTTP headers include various kinds of information, including the username and user group (if available).

User Authentication and Cloud Web Security

User identity can be used to apply policy in Cloud Web Security. User identity is also useful for Cloud Web Security reporting. User identity is not required to use Cloud Web Security. There are other methods to identify traffic for Cloud Web Security policy.

The ASA supports the following methods of determining the identity of a user, or of providing a default identity:

- AAA rules—When the ASA performs user authentication using a AAA rule, the username is retrieved from the AAA server or local database. Identity from AAA rules does not include group information. If configured, the default group is used. For information about configuring AAA rules, see Chapter 53, "Configuring AAA Rules for Network Access."
- IDFW—When the ASA uses IDFW with the Active Directory (AD), the username and group is retrieved from the AD agent when you activate a user and/or group by using an ACL in a feature such as an access rule or in your service policy, or by configuring the user identity monitor to download user identity information directly.

  For information about configuring IDFW, see Chapter 47, "Configuring the Identity Firewall."
- Default username and group—Without user authentication, the ASA uses an optional default username and/or group for all users that match a service policy rule for Cloud Web Security.

Authentication Keys

Each ASA must use an authentication key that you obtain from Cloud Web Security. The authentication key lets Cloud Web Security identify the company associated with web requests and ensures that the ASA is associated with valid customer.

You can use one of two types of authentication keys for your ASA: the company key or the group key.

- Company Authentication Key, page 75-3
- Group Authentication Key, page 75-3

Company Authentication Key

A Company authentication key can be used on multiple ASAs within the same company. This key simply enables the Cloud Web Security service for your ASAs. The administrator generates this key in ScanCenter (https://scancenter.scansafe.com/portal/admin/login.jsp); you have the opportunity to e-mail the key for later use. You cannot look up this key later in ScanCenter; only the last 4 digits are shown in ScanCenter. For more information, see the Cloud Web Security documentation: http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html.

Group Authentication Key

A Group authentication key is a special key unique to each ASA that performs two functions:

- Enables the Cloud Web Security service for one ASA.
- Identifies all traffic from the ASA so you can create ScanCenter policy per ASA.

For information about using the Group authentication key for policy, see the "ScanCenter Policy" section on page 75-4).

The administrator generates this key in ScanCenter (https://scancenter.scansafe.com/portal/admin/login.jsp); you have the opportunity to e-mail the key for later use. You cannot look up this key later in ScanCenter; only the last 4 digits are shown in ScanCenter. For more information, see the Cloud Web Security documentation: http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html.

ScanCenter Policy

In ScanCenter, traffic is matched against policy rules in order until a rule is matched. Cloud Web Security then applies the configured action for the rule. User traffic can match a policy rule in ScanCenter based on group association: a *directory group* or a *custom group*.

- Directory Groups, page 75-4
- Custom Groups, page 75-4
- How Groups and the Authentication Key Interoperate, page 75-5

Directory Groups

Directory groups define the group to which traffic belongs. The group, if present, is included in the HTTP header of the client request. The ASA includes the group in the HTTP header when you configure IDFW. If you do not use IDFW, you can configure a default group for traffic matching an ASA rule for Cloud Web Security inspection.

When you configure a directory group, you must enter the group name exactly.

- IDFW group names are sent in the following format:

  *domain-name**group-name*

  When the ASA learns the IDFW group name, the format on the ASA is *domain-name**group-name*. However, the ASA modifies the name to use only one backslash (\) to conform to typical ScanCenter notation.

- The default group name is sent in the following format:

  [*domain*\]*group-name*

  On the ASA, you need to configure the optional domain name to be followed by 2 backslashes (\\); however, the ASA modifies the name to use only one backslash (\) to conform to typical ScanCenter notation. For example, if you specify “Cisco\\Boulder1,” the ASA modifies the group name to be “Cisco\Boulder1” with only one backslash (\) when sending the group name to Cloud Web Security.

Custom Groups

Custom groups are defined using one or more of the following criteria:

- ScanCenter Group authentication key—You can generate a Group authentication key for a custom group. Then, if you identify this group key when you configure the ASA, all traffic from the ASA is tagged with the Group key.
- Source IP address—You can identify source IP addresses in the custom group. Note that the ASA service policy is based on source IP address, so you might want to configure any IP address-based policy on the ASA instead.
- Username—You can identify usernames in the custom group.
  - IDFW usernames are sent in the following format:

    *domain-name**username*
  - AAA usernames, when using RADIUS or TACACS+, are sent in the following format:

    LOCAL*username*
  - AAA usernames, when using LDAP, are sent in the following format:

    *domain-name**username*

– For the default username, it is sent in the following format:

[*domain-name*\]*username*

For example, if you configure the default username to be “Guest,” then the ASA sends “Guest.” If you configure the default username to be “Cisco\Guest,” then the ASA sends “Cisco\Guest.”

How Groups and the Authentication Key Interoperate

Unless you need the per-ASA policy that a custom group+group key provides, you will likely use a company key. Note that not all custom groups are associated with a group key. Non-keyed custom groups can be used to identify IP addresses or usernames, and can be used in your policy along with rules that use directory groups.

Even if you do want per-ASA policy and are using a group key, you can also use the matching capability provided by directory groups and non-keyed custom groups. In this case, you might want an ASA-based policy, with some exceptions based on group membership, IP address, or username. For example, if you want to exempt users in the America\Management group across all ASAs:

1. Add a directory group for America\Management.
2. Add an exempt rule for this group.
3. Add rules for each custom group+group key after the exempt rule to apply policy per-ASA.
4. Traffic from users in America\Management will match the exempt rule, while all other traffic will match the rule for the ASA from which it originated.

Many combinations of keys, groups, and policy rules are possible.

Cloud Web Security Actions

After applying the configured policies, Cloud Web Security either blocks, allows, or sends a warning about the user request:

- Allows—When Cloud Web Security allows the client request, it contacts the originally requested server and retrieves the data. It forwards the server response to the ASA, which then forwards it to the user.
- Blocks—When Cloud Web Security blocks the client request, it notifies the user that access has been blocked. It sends an HTTP 302 “Moved Temporarily” response that redirects the client application to a web page hosted by the Cloud Web Security proxy server showing the blocked error message. The ASA forwards the 302 response to the client.
- Warns—When the Cloud Web Security proxy server determines that a site may be in breach of the acceptable use policy, it displays a warning page about the site. You can choose to heed the warning and drop the request to connect, or you can click through the warning and proceed to the requested site.

You can also choose how the ASA handles web traffic when it cannot reach either the primary or backup Cloud Web Security proxy server. It can block or allow all web traffic. By default, it blocks web traffic.

Bypassing Scanning with Whitelists

If you use AAA rules or IDFW, you can configure the ASA so that web traffic from specific users or groups that otherwise match the service policy rule is not redirected to the Cloud Web Security proxy server for scanning. When you bypass Cloud Web Security scanning, the ASA retrieves the content

directly from the originally requested web server without contacting the proxy server. When it receives the response from the web server, it sends the data to the client. This process is called “whitelisting” traffic.

Although you can achieve the same results of exempting traffic based on user or group when you configure the class of traffic using ACLs to send to Cloud Web Security, you might find it more straightforward to use a whitelist instead. Note that the whitelist feature is only based on user and group, not on IP address.

IPv4 and IPv6 Support

Cloud Web Security currently supports only IPv4 addresses. If you use IPv6 internally, NAT 64 must be performed for any IPv6 flows that need to be sent to Cloud Web Security.

The following table shows the class map traffic that is supported by Cloud Web Security redirection:

Class Map Traffic	Cloud Web Security Inspection
From IPv4 to IPv4	Supported
From IPv6 to IPv4 (using NAT64)	Supported
From IPv4 to IPv6	Not Supported
From IPv6 to IPv6	Not Supported

Failover from Primary to Backup Proxy Server

When you subscribe to the Cisco Cloud Web Security service, you are assigned a primary Cloud Web Security proxy server and backup proxy server. These servers are routinely polled to check for their availability. If your ASA is unable to reach the Cloud Web Security proxy server (for example, if no SYN/ACK packets arrive from the proxy server), then the proxy server is polled through a TCP three-way handshake to check its availability. If the proxy server is unavailable after a configured number of retries (default is five), the server is declared as unreachable, and the backup proxy server becomes active.

The ASA automatically falls back to the primary Cloud Web Security proxy server from the backup server after continued polling shows that the primary server is active for two consecutive retry count periods. You can change this polling interval using the CLI or ASDM. For configuration information, see the “Configuring Communication with the Cloud Web Security Proxy Server” section on page 75-8.

Traffic Conditions Under Which Proxy Server Is Not Reachable	Server Timeout Calculation	Connection Timeout Result
High traffic	Client half open connection timeout + ASA TCP connection timeout	(30 + 30) = 60 seconds
Single connection failure	Client half open connection timeout + ((retry threshold - 1) x (ASA TCP connection timeout))	(30 + ((5-1) x (30)) = 150 seconds
Idle—No connections are passing	15 minutes + ((retry threshold) x (ASA TCP connection timeout))	900 + (5 x (30) = 1050 seconds

Licensing Requirements for Cisco Cloud Web Security

Model	License Requirement
All models	Strong Encryption (3DES/AES) License to encrypt traffic between the ASA and the Cloud Web Security server.

On the Cloud Web Security side, you must purchase a Cisco Cloud Web Security license and identify the number of users that the ASA handles. Then log into ScanCenter, and generate your authentication keys.

Prerequisites for Cloud Web Security

(Optional) User Authentication Prerequisites

To send user identity information to Cloud Web Security, configure one of the following on the ASA:

- AAA rules (username only)—See Chapter 53, "Configuring AAA Rules for Network Access."
- IDFW (username and group)—See Chapter 47, "Configuring the Identity Firewall."

(Optional) Fully Qualified Domain Name Prerequisites

If you use FQDNs in ACLs for your service policy rule, or for the Cloud Web Security server, you must configure a DNS server for the ASA according to the "Configuring the DNS Server" section on page 21-8.

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context modes.

In multiple context mode, the server configuration is allowed only in the system, and the service policy rule configuration is allowed only in the security contexts.

Each context can have its own authentication key, if desired.

Firewall Mode Guidelines

Supported in routed firewall mode only. Does not support transparent firewall mode.

IPv6 Guidelines

Does not support IPv6. See the "IPv4 and IPv6 Support" section on page 75-6.

Additional Guidelines

- Cloud Web Security is not supported with ASA clustering.
- Clientless SSL VPN is not supported with Cloud Web Security; be sure to exempt any clientless SSL VPN traffic from the ASA service policy for Cloud Web Security.

- When an interface to the Cloud Web Security proxy servers goes down, output from the **show scansafe server** command shows both servers up for approximately 15-25 minutes. This condition may occur because the polling mechanism is based on the active connection, and because that interface is down, it shows zero connection, and it takes the longest poll time approach.
- Cloud Web Security is not supported with the ASA CX module. If you configure both the ASA CX action and Cloud Web Security inspection for the same traffic, the ASA only performs the ASA CX action.
- Cloud Web Security inspection is compatibile with HTTP inspection for the same traffic. HTTP inspection is enabled by default as part of the default global policy.
- Cloud Web Security is not supported with extended PAT or any application that can potentially use the same source port and IP address for separate connections. For example, if two different connections (targeted to separate servers) use extended PAT, the ASA might reuse the same source IP and source port for both connection translations because they are differentiated by the separate destinations. When the ASA redirects these connections to the Cloud Web Security server, it replaces the destination with the Cloud Web Security server IP address and port (8080 by default). As a result, both connections now appear to belong to the same flow (same source IP/port and destination IP/port), and return traffic cannot be untranslated properly.
- The Default Inspection Traffic traffic class does not include the default ports for the Cloud Web Security inspection (80 and 443).

Default Settings

By default, Cisco Cloud Web Security is not enabled.

Configuring Cisco Cloud Web Security

- Configuring Communication with the Cloud Web Security Proxy Server, page 75-8
- (Multiple Context Mode) Allowing Cloud Web Security Per Security Context, page 75-9
- Configuring a Service Policy to Send Traffic to Cloud Web Security, page 75-10
- (Optional) Configuring Whitelisted Traffic, page 75-23
- Configuring the Cloud Web Security Policy, page 75-26

Configuring Communication with the Cloud Web Security Proxy Server

Guidelines

The public key is embedded in the ASA software, so there is no need for you to configure it.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Cloud Web Security**.

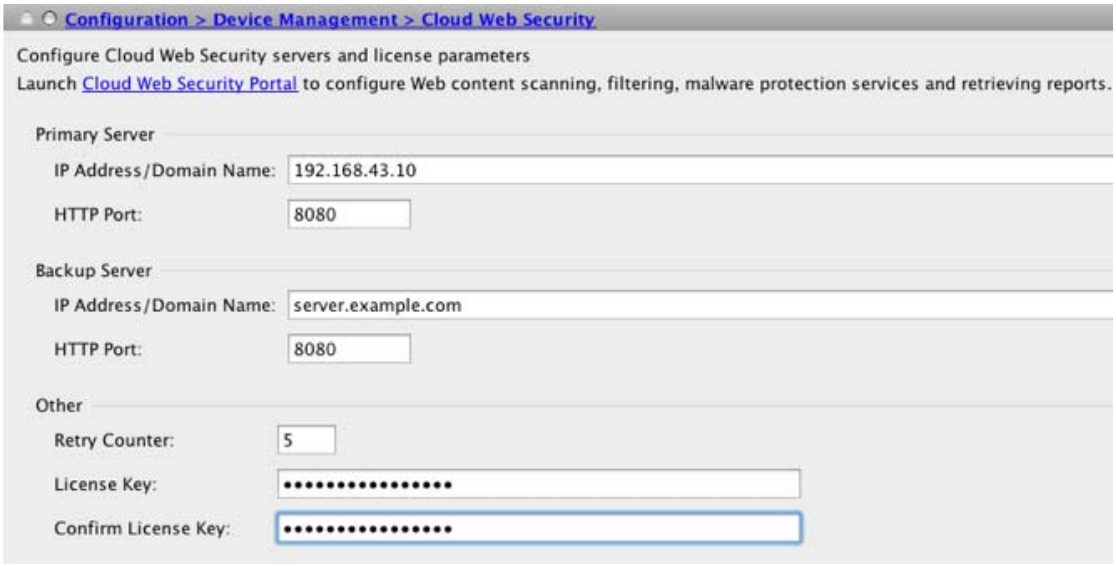

Step 2 In the Primary Server area, enter the following:

- IP Address/Domain Name—Enter the IPv4 address or FQDN of the primary server.
- HTTP Port—Enter the HTTP port of the primary server (port to which traffic must be redirected). By default the port is 8080; do not change this value unless directed to do so.

Step 3 In the Backup Server area, enter the following:

- IP Address/Domain Name—Enter the IPv4 address or FQDN of the backup server.
- HTTP Port—Enter the HTTP port of the backup server (port to which traffic must be redirected). By default the port is 8080. Valid values are from 1 to 65535.

Step 4 In the Other area, enter the following:

- Retry Counter—Enter the retry counter. Valid values are from 2 to 100.
- License Key—Enter the license (authentication) key, which is a 16-byte hexidecimal key. See the "Authentication Keys" section on page 75-3.
- Confirm License Key—Confirm the authentication key.

Step 5 Click **Apply**.

(Multiple Context Mode) Allowing Cloud Web Security Per Security Context

In multiple context mode, you must allow Cloud Web Security per context. See the "Configuring a Security Context" section on page 9-20.

Configuring a Service Policy to Send Traffic to Cloud Web Security

Your service policy consists of multiple service policy rules, applied globally, or applied to each interface. Each service policy rule can either send traffic to Cloud Web Security (Match) or exempt traffic from Cloud Web Security (Do Not Match). Create rules for traffic destined for the Internet. The order of these rules is important. When the ASA decides whether to forward or exempt a packet, the ASA tests the packet with each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create a rule at the beginning of a policy that explicitly Matches all traffic, no further statements are ever checked. You can reorder the rules as needed after you add them.

See Chapter 44, "Configuring a Service Policy," for more information about service policy rules.

Prerequisites

(Optional) If you need to use a whitelist to exempt some traffic from being sent to Cloud Web Security, first create the whitelist according to the "(Optional) Configuring Whitelisted Traffic" section on page 75-23 so you can refer to the whitelist in your service policy rule.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule** to add a service policy rule.

Step 2 On the Service Policy dialog box, you can configure Cloud Web Security as part of a new service policy, or you can edit an existing service policy. Click **Next**.

Step 3 On the Traffic Classification Criteria dialog box, name the traffic class (or accept the default name), keep the **Create a new traffic class** option selected, and click **Source and Destination IP address (Uses ACL)**, then click **Next**.

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

The Traffic Match - Source and Destination dialog box appears.

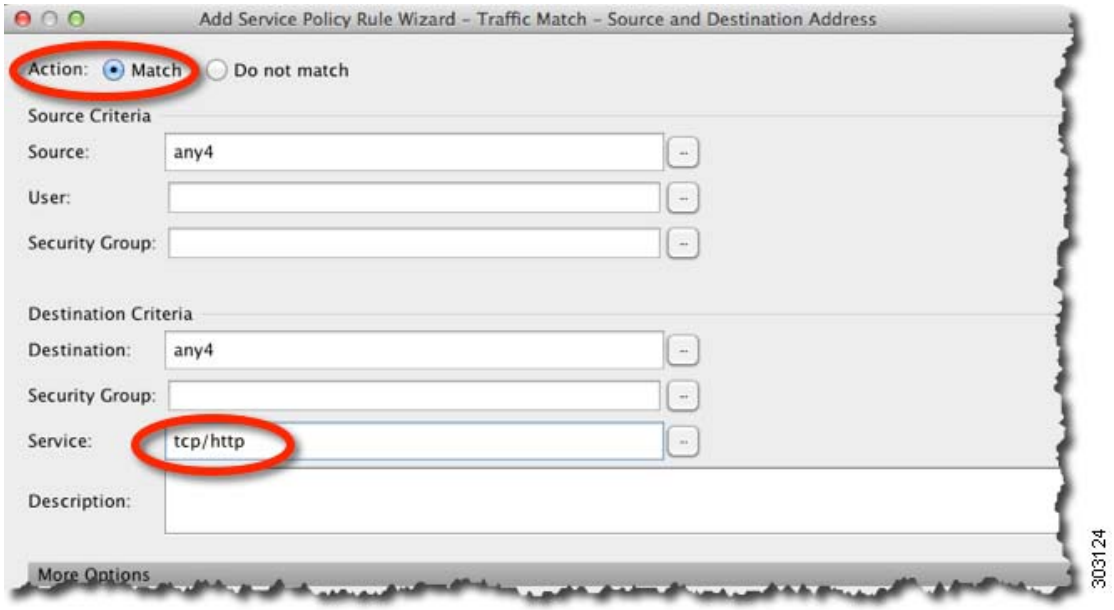

a. Click **Match** or **Do Not Match**.

   Match specifies that traffic matching the source and destination is sent to Cloud Web Security. **Do Not Match** exempts matching traffic from Cloud Web Security. You can later add additional rules to match or not match other traffic.

   When creating your rules, consider how you can match appropriate traffic that is destined for the Internet, but not match traffic that is destined for other internal networks. For example, to prevent inside traffic from being sent to Cloud Web Security when the destination is an internal server on the DMZ, be sure to add a deny ACE to the ACL that exempts traffic to the DMZ.

b. In the Source Criteria area, enter or browse for a Source IP address or network object, an optional IDFW Username or group, and an optional TrustSec Security Group.

c. In the Destination Criteria area, enter or browse for a Destination IP address or network object, and an optional TrustSec Security Group.

   FQDN network objects might be useful in matching or exempting traffic to specific servers.

d. In the Service field, enter **http** or **https**, and click **Next**.

Note Cloud Web Security only operates on HTTP and HTTPS traffic. Each type of traffic is treated separately by the ASA. Therefore, you need to create HTTP-only rules and HTTPS-only rules.

The Rule Actions dialog box appears.

Step 4 On the Protocol Inspection tab, check the **Cloud Web Security** check box.

Step 5 Click **Configure** to set the traffic action (fail open or fail close) and add the inspection policy map.

The inspection policy map configures essential parameters for the rule and also optionally identifies the whitelist. An inspection policy map is required for each class of traffic that you want to send to Cloud Web Security. You can also pre-configure inspection policy maps from the Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security pane.

The Select Cloud Web Security Inspect Map dialog box appears.

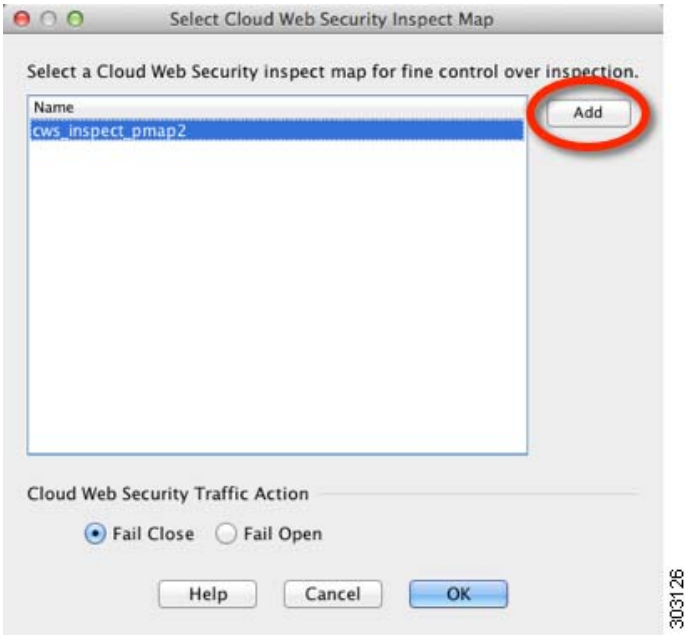

a. For the Cloud Web Security Traffic Action, choose one:
   - **Fail Close**—Drops all traffic if the Cloud Web Security servers are unavailable.
   - **Fail Open**—Allows traffic to pass through the ASA if the Cloud Web Security servers are unavailable.

b. Choose an existing inspection policy map, or add one using the **Add** button.

c. Click **Add** to add a new inspection policy map.

   The Add Cloud Web Security Inspect Map dialog box appears.

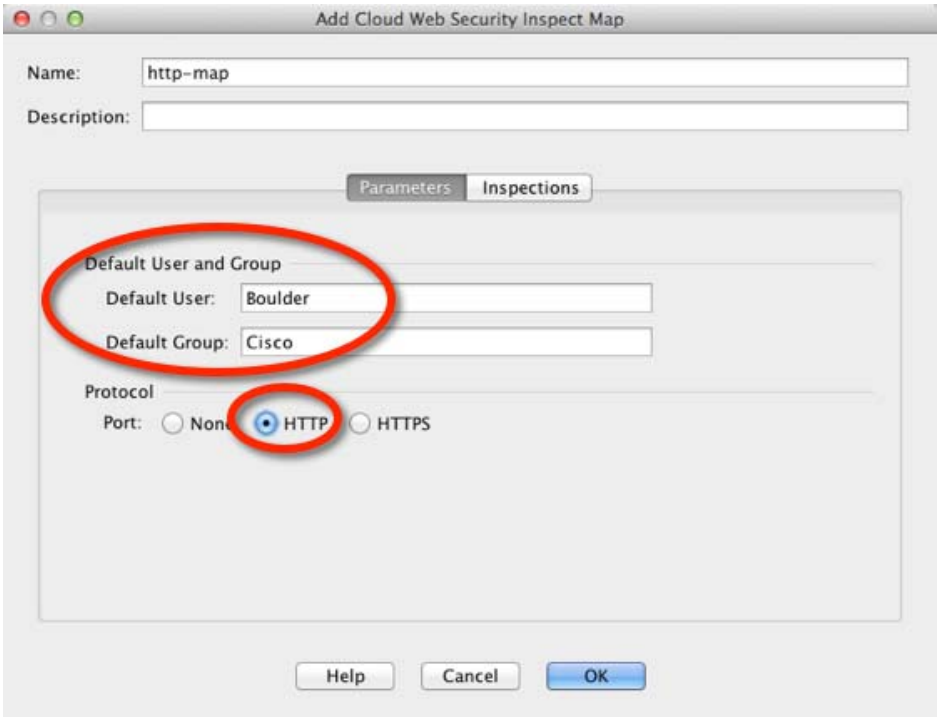

d. In the Name field, specify a name for the inspection policy map, up to 40 characters in length.

e. (Optional) Enter a description.

f. (Optional) On the Parameters tab, specify a Default User and/or a Default Group. If the ASA cannot determine the identity of the user coming into the ASA, then the default user and/or group is applied.

g. For the Protocol, click **HTTP** or **HTTPS**, to match the service you set in Step 3d. Cloud Web Security treats each type of traffic separately.

h. (Optional) To identify a whitelist, click the **Inspections** tab.

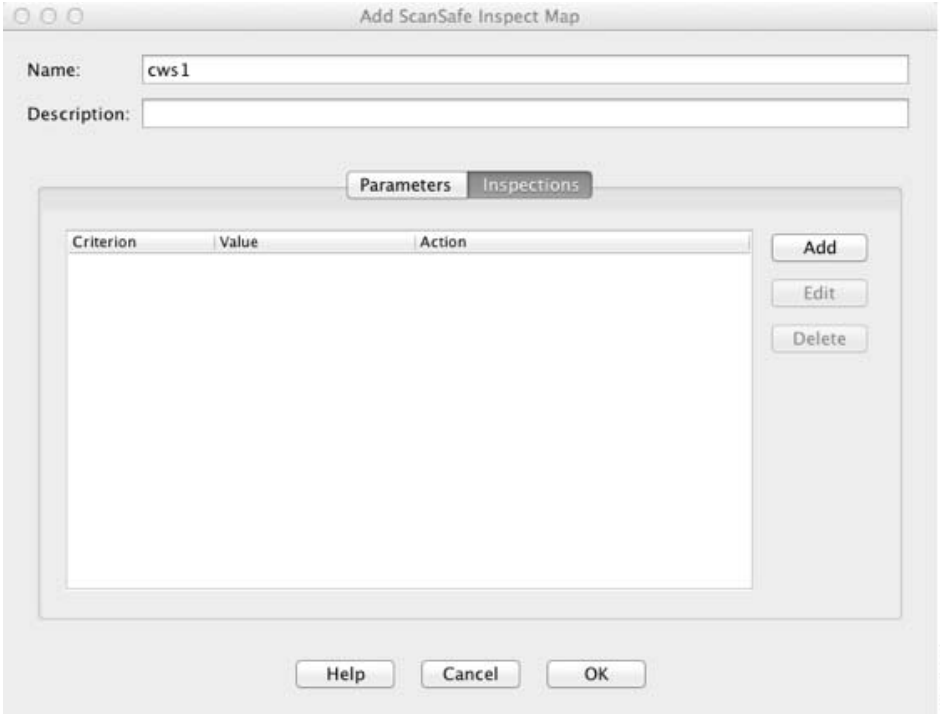

- Click **Add** to choose the inspection class map you created in the "(Optional) Configuring Whitelisted Traffic" section on page 75-23.

  The Add Cloud Web Security Match Criterion dialog box appears.
- From the Cloud Web Security Traffic Class drop-down menu, choose an inspection class map.

  To add or edit a class map, click **Manage**.
- For the Action, click **Whitelist**.
- Click **OK** to add the whitelist to the policy map.
- Click **OK**.

Step 6 Click **Finish**. The rule is added to the Service Policy Rules table.

Step 7 To add additional sub-rules (ACEs) for this traffic class, to match or exempt additional traffic:

a. Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**.

b. Choose the same service policy as Step 2. Click **Next**.

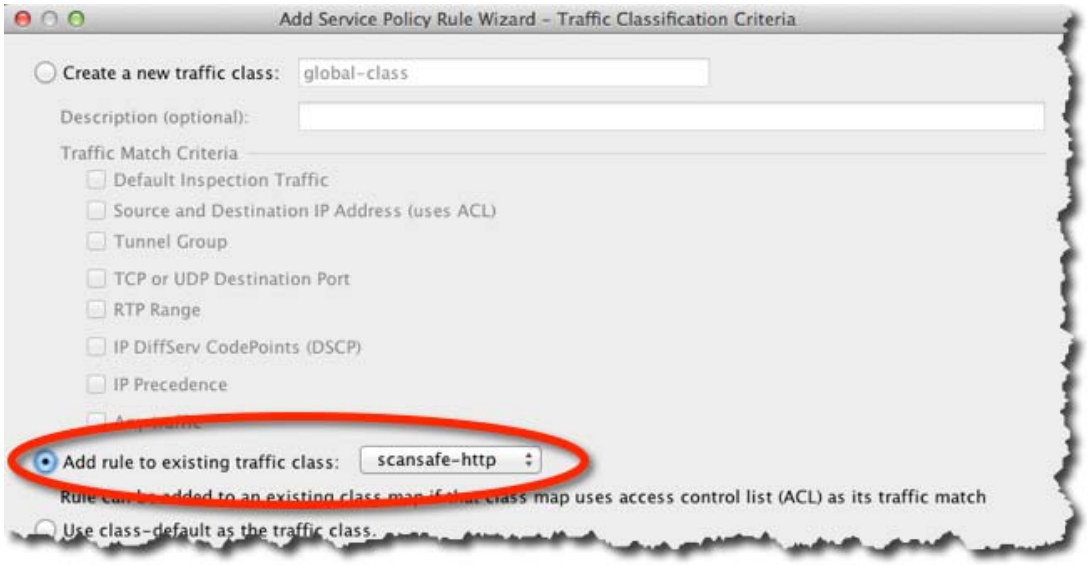

c. On the Traffic Classification Criteria dialog box, choose **Add Rule to Existing Traffic Class**, and choose the name you created in Step 3. Click **Next**.

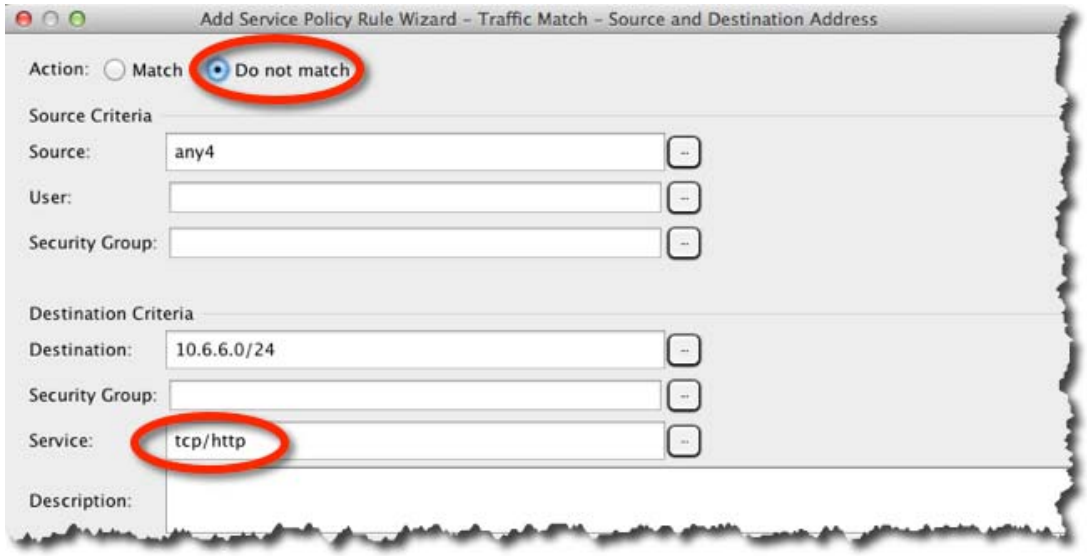

d. In the Traffic Match - Source and Destination dialog box, choose **Match** to add inspect additional traffic, or **Do Not Match** to exempt traffic from Cloud Web Security inspections. Be sure to set the service to match the previous rules in this class (HTTP or HTTPS); you cannot mix HTTP and HTTPS in the same traffic class for Cloud Web Security. Click **Next**.

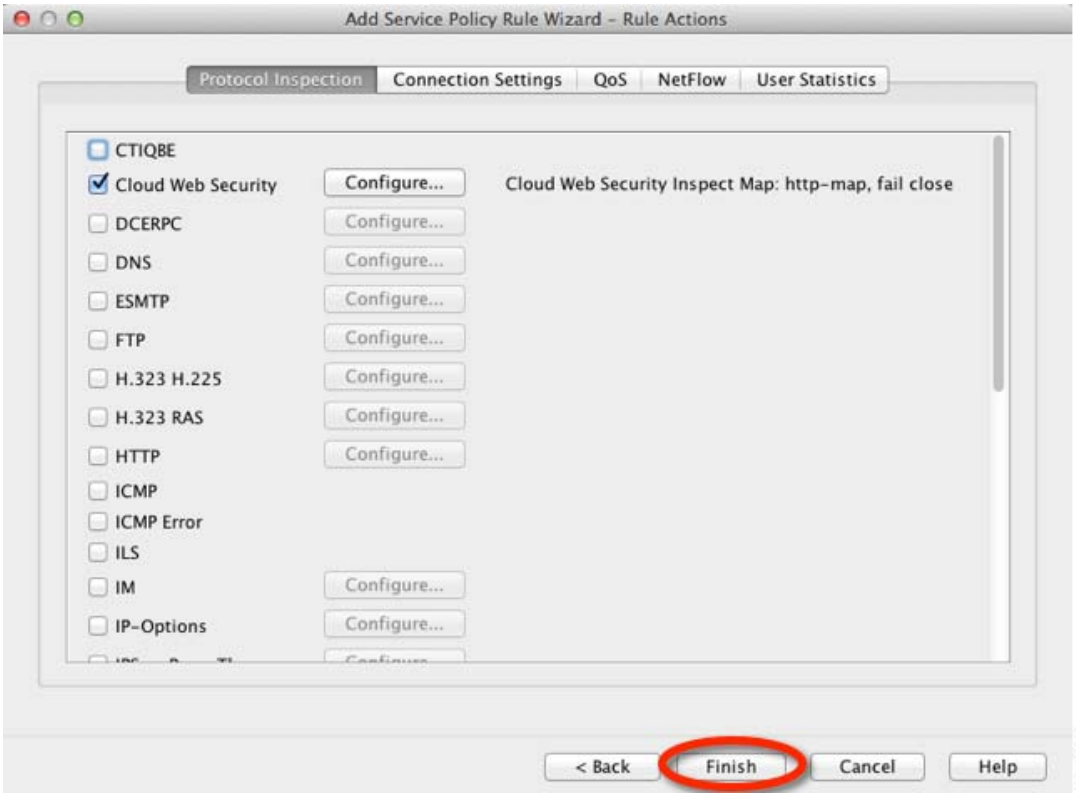

e. On the Rule Actions dialog box, do not make any changes; click **Finish**. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs, so the previously-specified actions are inherited.

Step 8 Repeat this entire procedure to create an additional traffic class, for example for HTTPS traffic. You can create as many rules and sub-rules as needed.

Step 9 Arrange the order of Cloud Web Security rules and sub-rules on the Service Policy Rules pane. See the "Managing the Order of Service Policy Rules" section on page 44-16 for information about changing the order of ACEs.

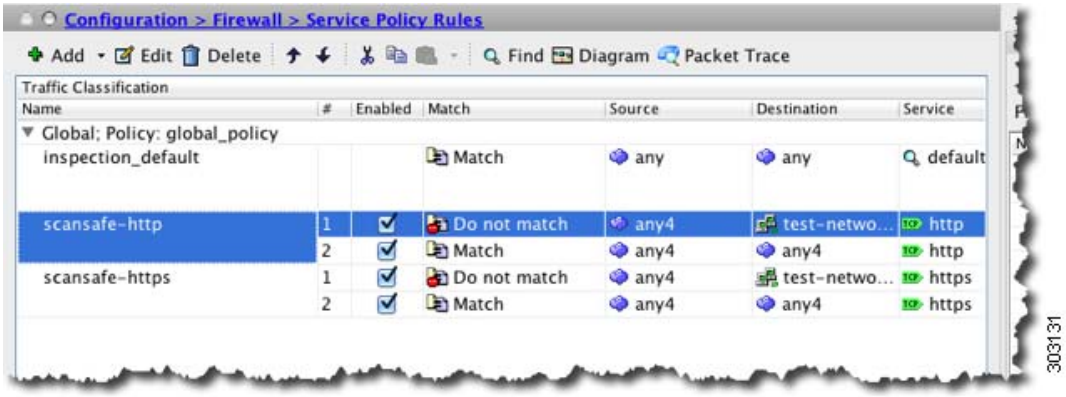

Step 10 Click **Apply**.

Examples

The following example exempts all IPv4 HTTP and HTTPS traffic going to the 10.6.6.0/24 (test_network), and sends all other HTTPS and HTTPS traffic to Cloud Web Security, and applies this service policy rule to all interfaces as part of the existing global policy. If the Cloud Web Security server is unreachable, the ASA drops all matching traffic (fail close). If a user is not have user identity information, the default user Boulder and group Cisco is used.

Step 1 Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**. Add this rule to the default global_policy:

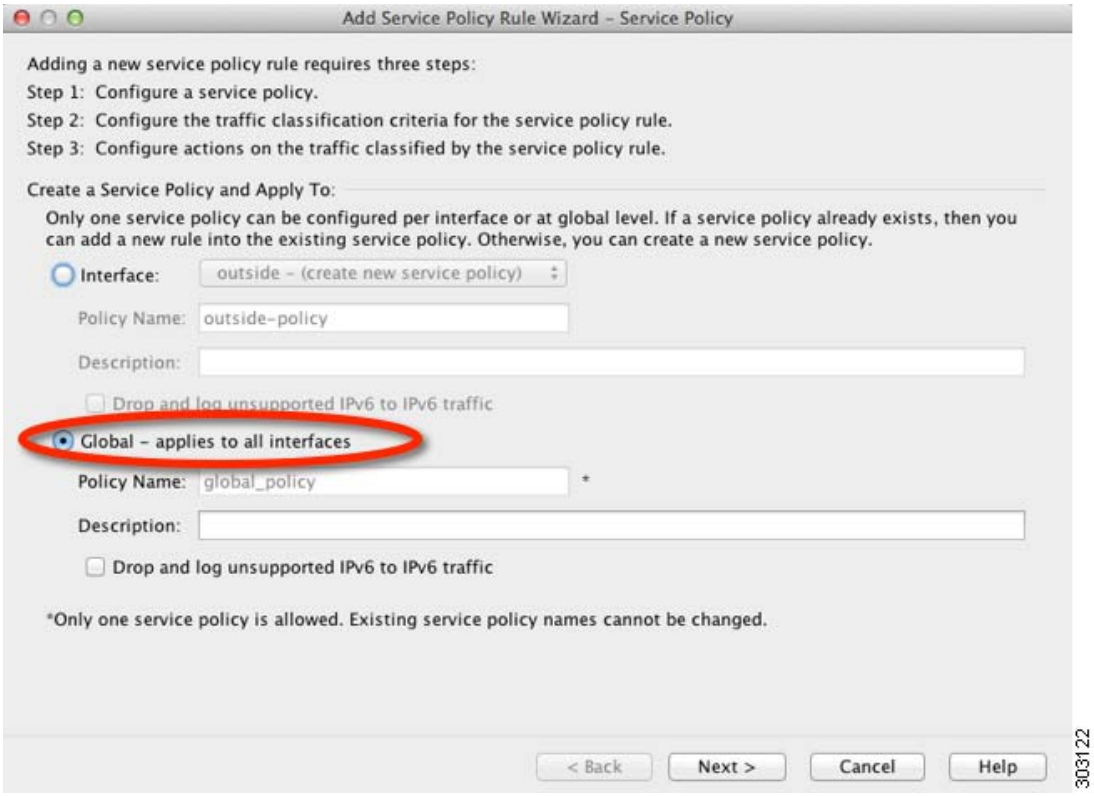

Step 2 Add a new traffic class called “scansafe-http,” and specify an ACL for traffic matching:

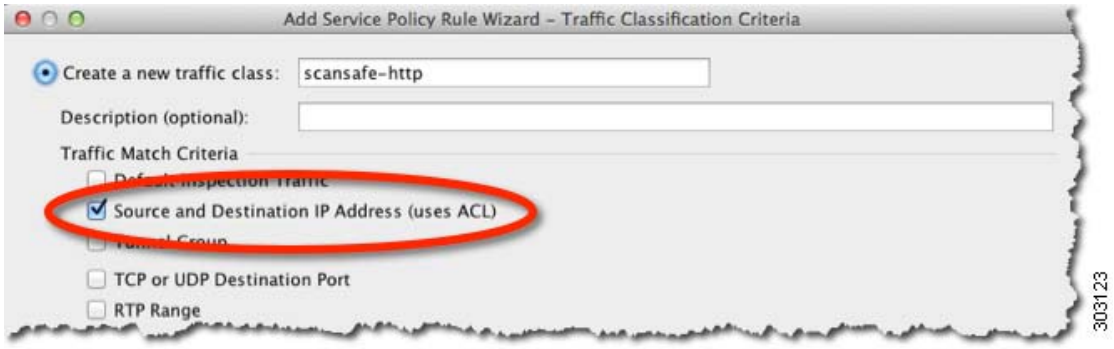

Step 3 Choose **Match**, and specify **any4** for the Source and Destination. Specify **tcp/http** for the Service.

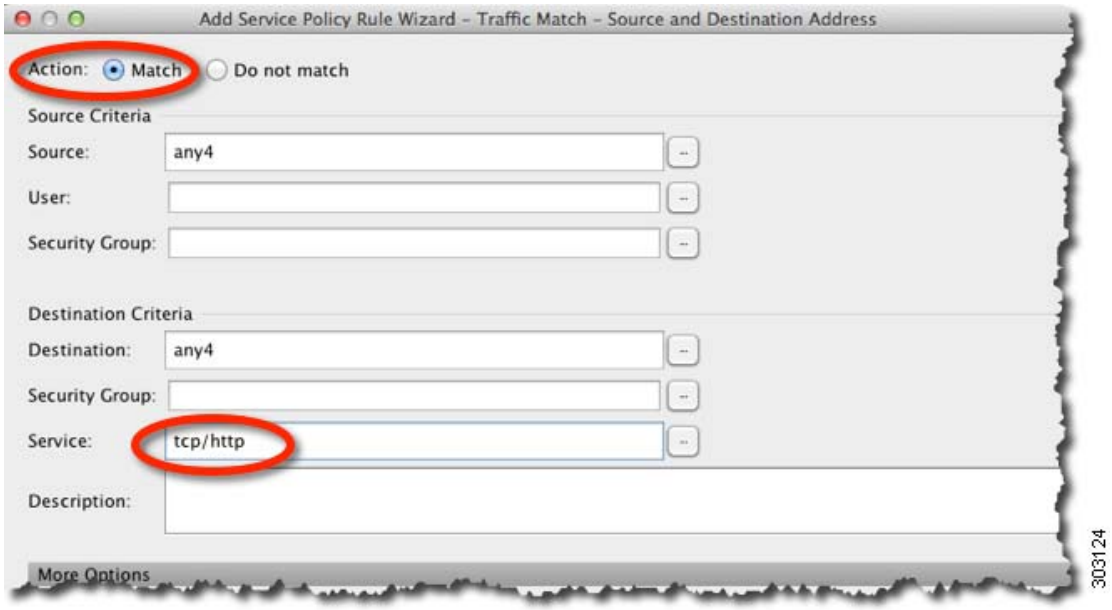

Step 4 Check **Cloud Web Security** and click **Configure**.

Step 5 Accept the default Fail Close action, and click **Add**.

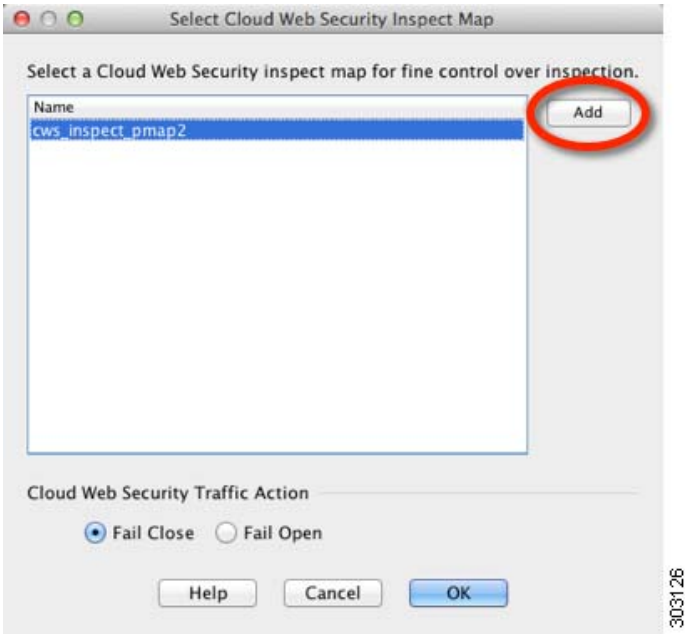

Step 6 Name the inspection policy map "http-map," set the Default User to Boulder and the default group to Cisco. Choose **HTTP**.

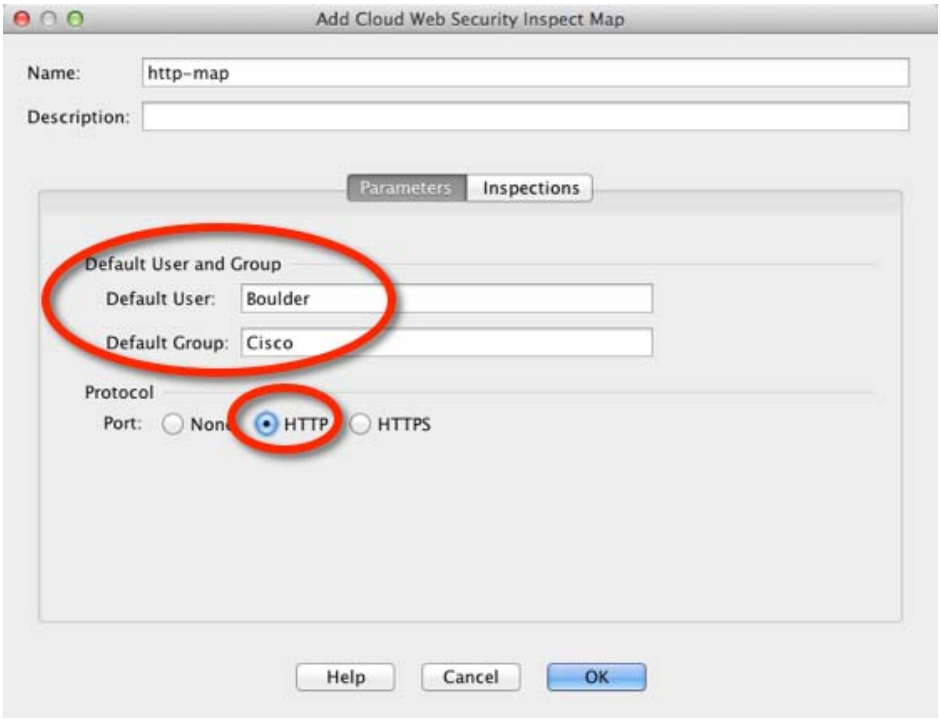

Step 7 Click **OK**, **OK**, and then **Finish**. The rule is added to the Service Policy Rules table.

Step 8 Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**. Add the new rule to the default global_policy:

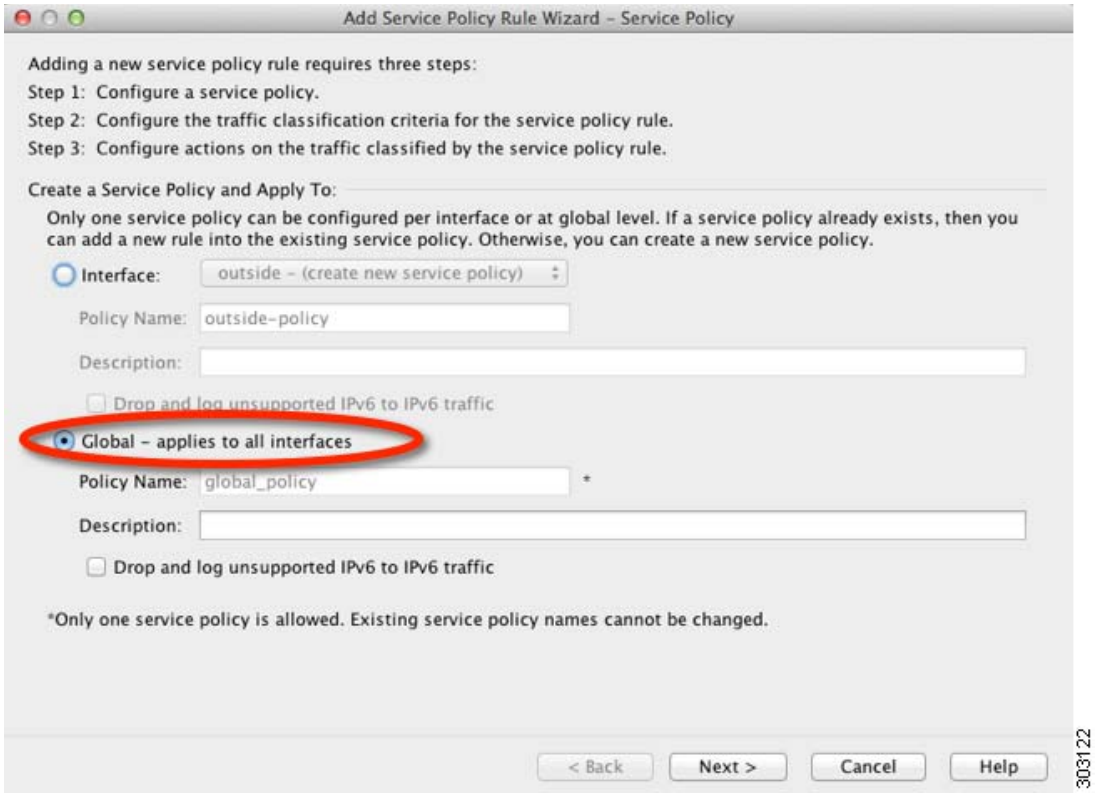

Step 9 Click **Add rule to existing traffic class**, and choose **scansafe-http**.

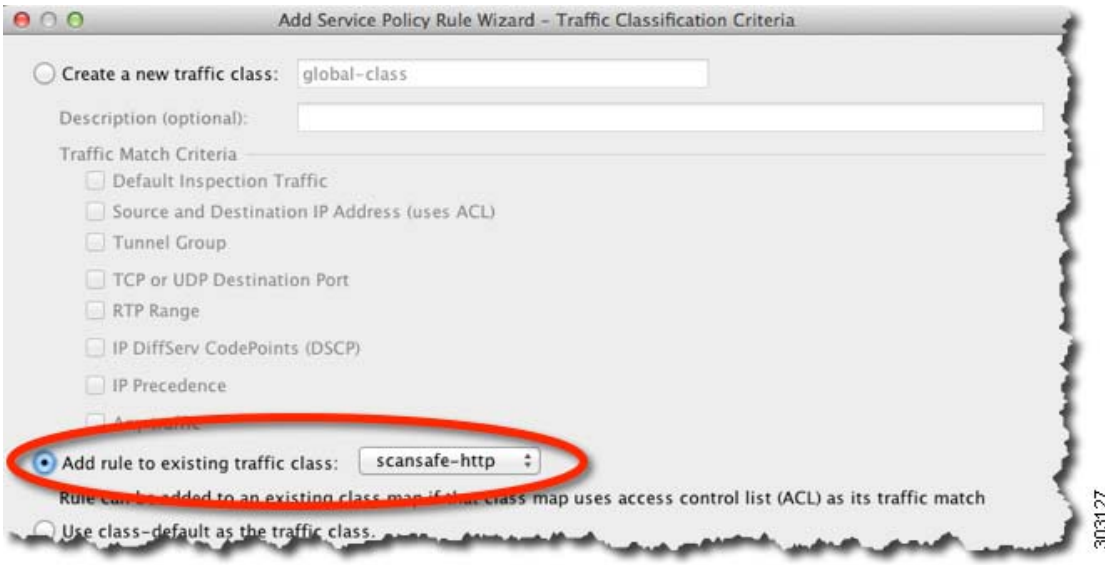

Step 10 Choose **Do not match**, set **any4** as the Source, and **10.6.6.0/24** as the Destination. Set the Service to **tcp/http**.

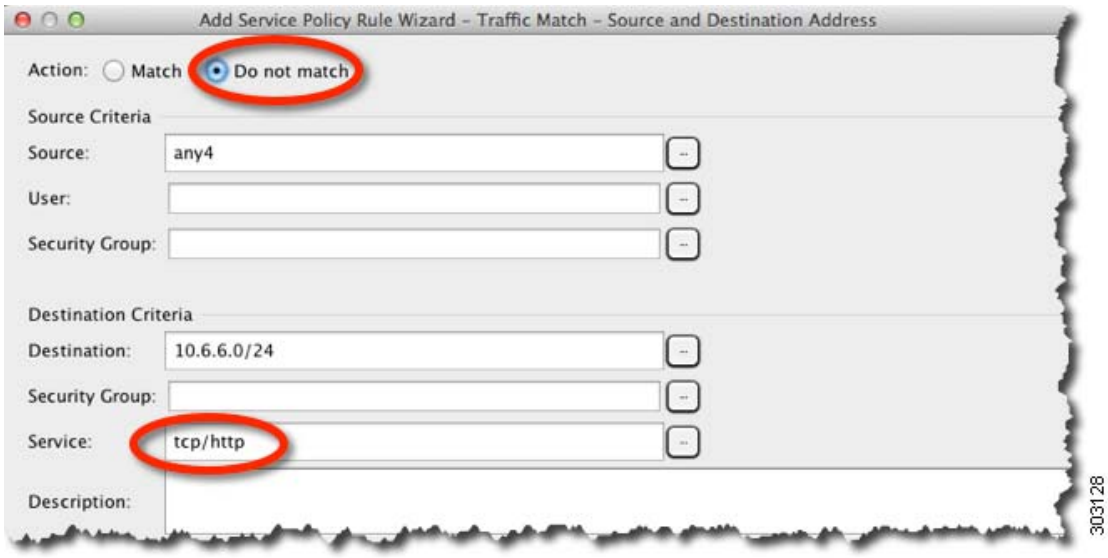

Step 11 Click **Finish**.

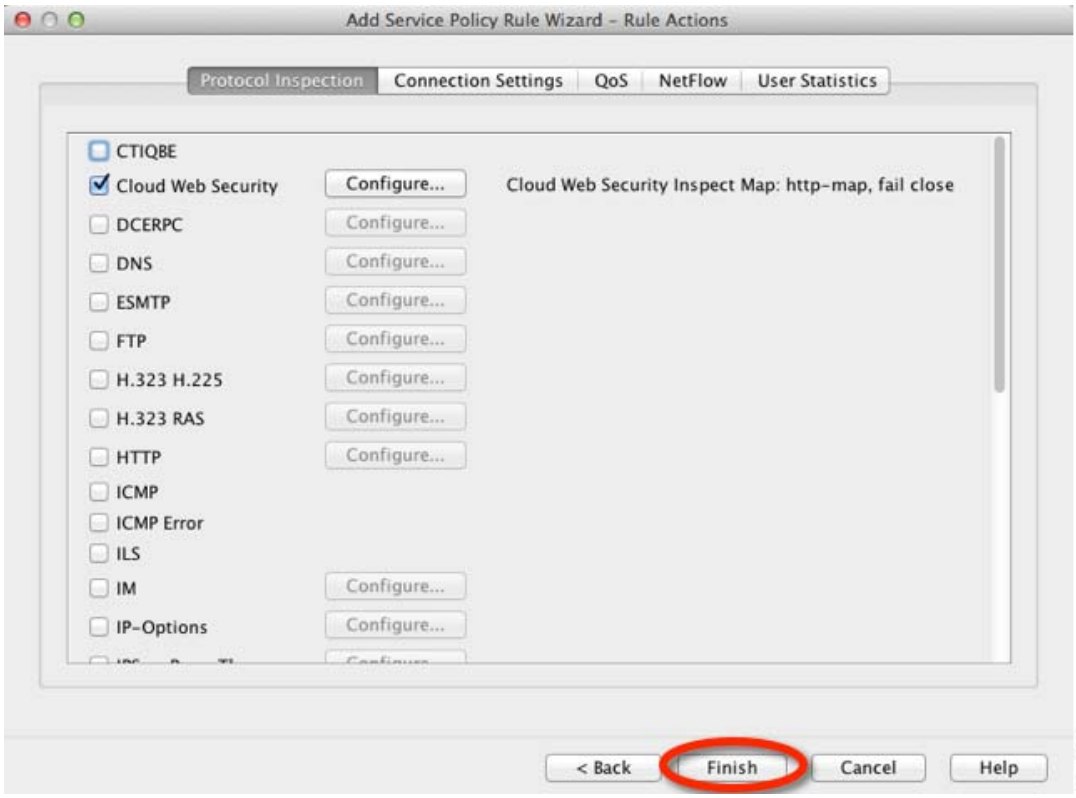

Step 12 Reorder the rules so the Do not match rule is above the Match rule.

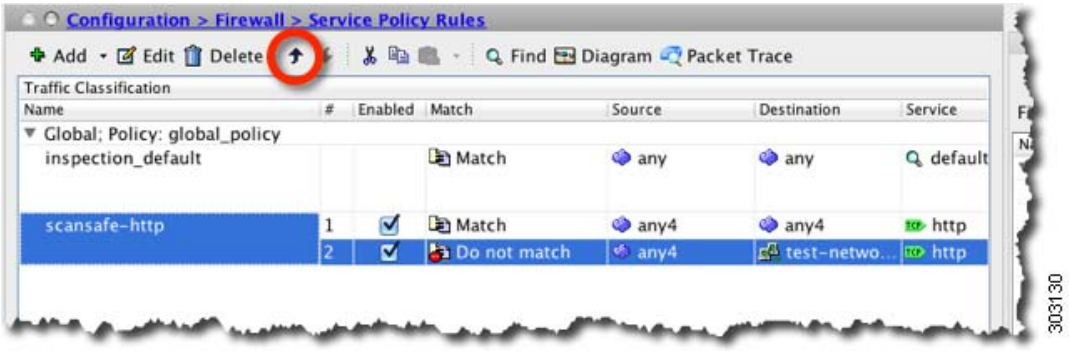

User traffic is compared to these rules in order; if this Match rule is first in the list, then all traffic, including traffic to test_network, will match only that rule and the Do not match rule will never be hit. If you move the Do not match rule above the Match rule, then traffic to test_network will match the Do not match rule, and all other traffic will match the Match rule.

Step 13 Repeat the above steps with the following changes: add a new traffic class called "scansafe-https," and choose **HTTPS** for the inspection policy map.

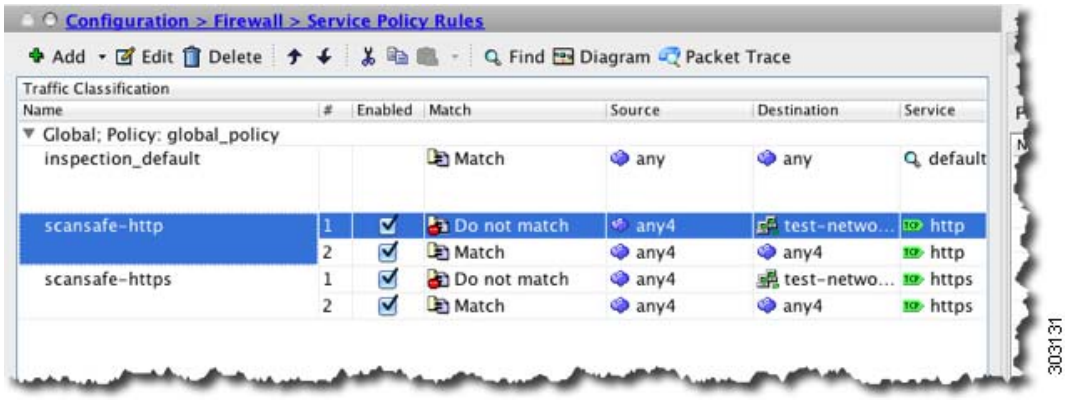

Step 14 Click **Apply**.

(Optional) Configuring Whitelisted Traffic

If you use user authentication, you can exempt some traffic from being filtered by Cloud Web Security based on the username and/or groupname. When you configure your Cloud Web Security service policy rule, you can reference the whitelisting inspection class map. Both IDFW and AAA user credentials can be used with this feature.

Although you can achieve the same results of exempting traffic based on user or group when you configure the service policy rule, you might find it more straightforward to use a whitelist instead. Note that the whitelist feature is only based on user and group, not on IP address.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Objects > Class Maps > Cloud Web Security**.

Step 2 Click **Add** to create a new class map.

The Add Cloud Web Security Traffic Class Map screen appears.

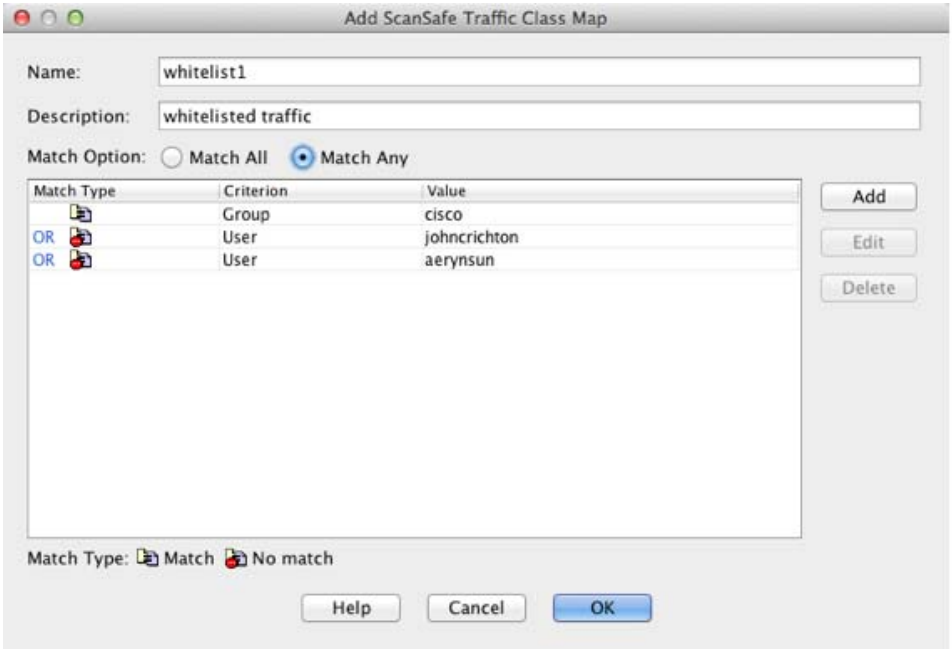

Step 3 In the Name field, enter the name of the new class map (40 characters or less).

Step 4 In the Description field, provide a description for the class map (200 characters or less).

Step 5 Choose the Match Option for the criteria you define when you click ADD:

- Match All—Specifies that traffic must match all criteria to match the class map.
- Match Any—Specifies that the traffic matches the class map if it matches at least one of the criteria.

Step 6 Click **Add**.

The Add Cloud Web Security Match Criterion Window appears.

Step 7 Choose the Match Type:

- Match—Specifies the user and/or group that you want to whitelist.
- No Match—Specifies the user and/or group that you do *not* want to whitelist; for example, if you whitelist the group “cisco,” but you want to scan traffic from users “johncrichton” and “aerynsun,” you can specify No Match for those users.

Step 8 Choose the Match Criterion:

- User—Specifies the user.
- Group—Specifies the group.
- User and Group—Specifies a user and group.

Step 9 Click **OK**.

Step 10 Continue to add match criteria as desired.

Step 11 Click **OK** to add the class map.

Step 12 Click **Apply**.

Step 13 Use the whitelist in the Cloud Web Security policy according to the "Configuring a Service Policy to Send Traffic to Cloud Web Security" section on page 75-10.

(Optional) Configuring the User Identity Monitor

When you use IDFW, the ASA only downloads user identity information from the AD server for users and groups included in active ACLs; the ACL must be used in a feature such as an access rule, AAA rule, service policy rule, or other feature to be considered active. Because Cloud Web Security can base its policy on user identity, you may need to download groups that are not part of an active ACL to get full IDFW coverage for all your users. For example, although you can configure your Cloud Web Security service policy rule to use an ACL with users and groups, thus activating any relevant groups, it is not required; you could use an ACL based entirely on IP addresses.The user identity monitor feature lets you download group information directly from the AD agent.

Restrictions

The ASA can only monitor a maximum of 512 groups, including those configured for the user identity monitor and those monitored through active ACLs.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Identity Options**, and scroll to the Cloud Web Security Configuration section.

Step 2 Click **Add**.

The Add Monitor User dialog box appears.

Step 3 To add a domain, click **Manage**, and then click **Add**. You can only monitor groups for domains you have pre-defined on the ASA.

The Configure Identity Domains dialog box appears. For detailed information about adding domains, see the "Configuring Identity Options" section on page 47-12.

Step 4 When you are finished adding domains, click **OK**.

Step 5 You can either type in a group name, or you can search for groups on the AD agent per domain.

- To type in a group name directly, enter the name in the bottom field in the following format, and click **OK**:

  *domain-name**group*

- To search for a group on the AD agent:

a. Choose the domain from the Domain drop-down list.

b. In the Find field, enter a text string to match group names, and click **Find**.

The ASA downloads names from the AD agent for the specified domain.

c. Double-click the name you want to monitor; it is added to the bottom field.

d. Click **OK**.

Repeat for additional groups.

Step 6 After you add the groups you want to monitor, click **Apply**.

Configuring the Cloud Web Security Policy

After you configure the ASA service policy rules, launch the ScanCenter Portal to configure Web content scanning, filtering, malware protection services, and reports.

Detailed Steps

Go to: https://scancenter.scansafe.com/portal/admin/login.jsp.

For more information, see the Cisco ScanSafe Cloud Web Security Configuration Guides:

http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html

Monitoring Cloud Web Security

Command	Purpose
Monitoring > Properties > Cloud Web Security	Shows the status of the server, whether it is the current active server, the backup server, or unreachable. Shows total and current HTTP(S) connections. In multiple context mode, statistics are only shown within a context.
See the following URL: http://Whoami.scansafe.net	From a client, access this web site to determine if your traffic is going to the Cloud Web Security server.

Related Documents

Related Documents	URL
Cisco ScanSafe Cloud Web Security Configuration Guides	http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html

Feature History for Cisco Cloud Web Security

Table 75-1 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 75-1 Feature History for Cloud Web Security

Feature Name	Platform Releases	Feature Information
Cloud Web Security	9.0(1)	This feature was introduced. Cisco Cloud Web Security provides content scanning and other malware protection service for web traffic. It can also redirect and report about web traffic based on user identity. We introduced or modified the following screens: Configuration > Device Management > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit > Manage Cloud Web Security Class Maps Configuration > Firewall > Identity Options Configuration > Firewall > Service Policy Rules Monitoring > Properties > Cloud Web Security

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 76-1
- Licensing Requirements for the Botnet Traffic Filter, page 76-6
- Guidelines and Limitations, page 76-6
- Default Settings, page 76-6
- Configuring the Botnet Traffic Filter, page 76-6
- Monitoring the Botnet Traffic Filter, page 76-14
- Where to Go Next, page 76-16
- Feature History for the Botnet Traffic Filter, page 76-16

Information About the Botnet Traffic Filter

This section includes information about the Botnet Traffic Filter and includes the following topics:

- Botnet Traffic Filter Address Types, page 76-2
- Botnet Traffic Filter Actions for Known Addresses, page 76-2
- Botnet Traffic Filter Databases, page 76-2

- How the Botnet Traffic Filter Works, page 76-5

Botnet Traffic Filter Address Types

Addresses monitored by the Botnet Traffic Filter include:

- Known malware addresses—These addresses are on the blacklist identified by the dynamic database and the static blacklist.
- Known allowed addresses—These addresses are on the whitelist. The whitelist is useful when an address is blacklisted by the dynamic database and also identified by the static whitelist.
- Ambiguous addresses—These addresses are associated with multiple domain names, but not all of these domain names are on the blacklist. These addresses are on the *greylist*.
- Unlisted addresses—These addresses are unknown, and not included on any list.

Botnet Traffic Filter Actions for Known Addresses

You can configure the Botnet Traffic Filter to log suspicious activity, and you can optionally configure it to block suspicious traffic automatically.

Unlisted addresses do not generate any syslog messages, but addresses on the blacklist, whitelist, and greylist generate syslog messages differentiated by type. See the "Botnet Traffic Filter Syslog Messaging" section on page 76-14 for more information.

Botnet Traffic Filter Databases

The Botnet Traffic Filter uses two databases for known addresses. You can use both databases together, or you can disable use of the dynamic database and use the static database alone. This section includes the following topics:

- Information About the Dynamic Database, page 76-2
- Information About the Static Database, page 76-3
- Information About the DNS Reverse Lookup Cache and DNS Host Cache, page 76-4

Information About the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses.

How the ASA Uses the Dynamic Database

The ASA uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the *DNS reverse lookup cache*.
2. When the infected host starts a connection to the IP address of the malware site, then the ASA sends a syslog message informing you of the suspicious activity and optionally drops the traffic if you configured the ASA to do so.

3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs or drops any traffic to that IP address without having to inspect DNS requests.

Database Files

The database files are stored in running memory; they are not stored in flash memory. If you need to delete the database, use theConfiguration > Firewall > Botnet Traffic Filter > Botnet Database pane Purge Botnet Database button instead. Be sure to first disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box in the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area.

Note To use the database, be sure to configure a domain name server for the ASA 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 ASA looks inside the DNS packets for the domain name and associated IP address.

Database Traffic Types

The dynamic database includes the following types of addresses:

- Ads—These are advertising networks that deliver banner ads, interstitials, rich media ads, pop-ups, and pop-unders for websites, spyware and adware. Some of these networks send ad-oriented HTML emails and email verification services.
- Data Tracking—These are sources associated with companies and websites that offer data tracking and metrics services to websites and other online entities. Some of these also run small advertising networks.
- Spyware—These are sources that distribute spyware, adware, greyware, and other potentially unwanted advertising software. Some of these also run exploits to install such software.
- Malware—These are sources that use various exploits to deliver adware, spyware and other malware to victim computers. Some of these are associated with rogue online vendors and distributors of dialers which deceptively call premium-rate phone numbers.
- Adult—These are sources associated with adult networks/services offering web hosting for adult content, advertising, content aggregation, registration & billing, and age verification. These may be tied to distribution of adware, spyware, and dialers.
- Bot and Threat Networks—These are rogue systems that control infected computers. They are either systems hosted on threat networks or systems that are part of the botnet itself.

Information About the Static Database

You can manually enter domain names or IP addresses (host or subnet) that you want to tag as bad names in a blacklist. Static blacklist entries are always designated with a Very High threat level. You can also enter names or IP addresses in a whitelist, so that names or addresses that appear on both the *dynamic* blacklist and the whitelist are identified only as whitelist addresses in syslog messages and reports. Note that you see syslog messages for whitelisted addresses even if the address is not also in the dynamic blacklist.

When you add a domain name to the static database, the ASA waits 1 minute, and then sends a DNS request for that domain name and adds the domain name/IP address pairing to the *DNS host cache*. (This action is a background process, and does not affect your ability to continue configuring the ASA). We

recommend also enabling DNS packet inspection with Botnet Traffic Filter snooping. The ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following circumstances:

- The ASA DNS server is unavailable.
- A connection is initiated during the 1 minute waiting period before the ASA sends the regular DNS request.

If DNS snooping is used, when an infected host sends a DNS request for a name on the static database, the ASA looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

If you do not enable Botnet Traffic Filter snooping, and one of the above circumstances occurs, then that traffic will not be monitored by the Botnet Traffic Filter.

Information About the DNS Reverse Lookup Cache and DNS Host Cache

When you use the dynamic database with DNS snooping, entries are added to the DNS reverse lookup cache. If you use the static database, entries are added to the DNS host cache (see the "Information About the Static Database" section on page 76-3 about using the static database with DNS snooping and the DNS reverse lookup cache).

Entries in the DNS reverse lookup cache and the DNS host cache have a time to live (TTL) value provided by the DNS server. The largest TTL value allowed is 1 day (24 hours); if the DNS server provides a larger TTL, it is truncated to 1 day maximum.

For the DNS reverse lookup cache, after an entry times out, the ASA renews the entry when an infected host initiates a connection to a known address, and DNS snooping occurs.

For the DNS host cache, after an entry times out, the ASA periodically requests a refresh for the entry.

For the DNS host cache, the maximum number of blacklist entries and whitelist entries is 1000 each.

Table 76-1 lists the maximum number of entries in the DNS reverse lookup cache per model.

Table 76-1 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 76-1 shows how the Botnet Traffic Filter works with the dynamic database plus DNS inspection with Botnet Traffic Filter snooping.

Figure 76-1 How the Botnet Traffic Filter Works with the Dynamic Database

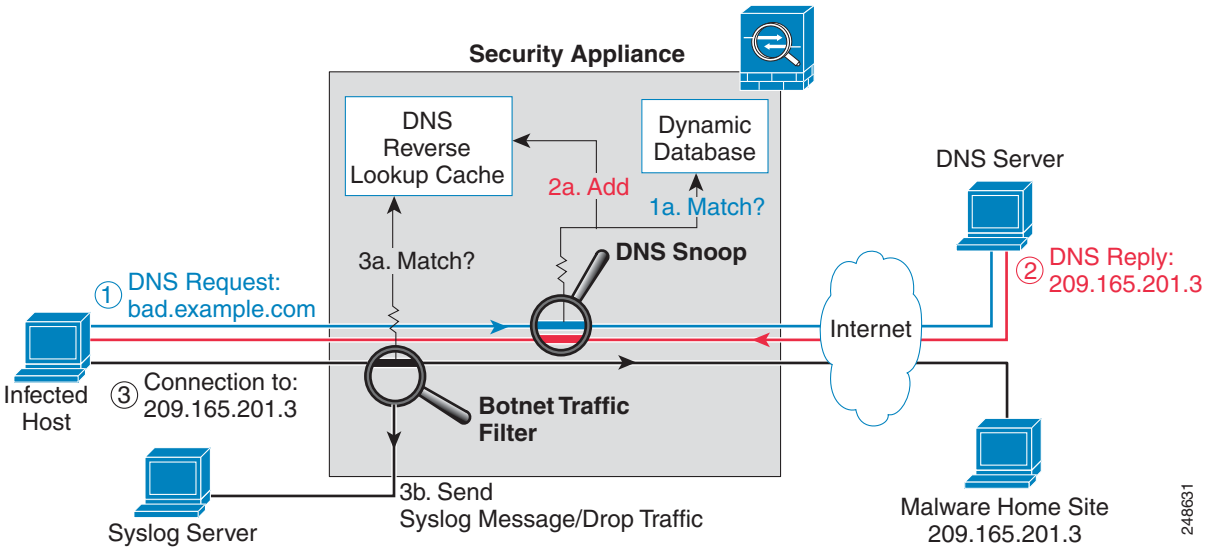

Figure 76-2 shows how the Botnet Traffic Filter works with the static database.

Figure 76-2 How the Botnet Traffic Filter Works with the Static Database

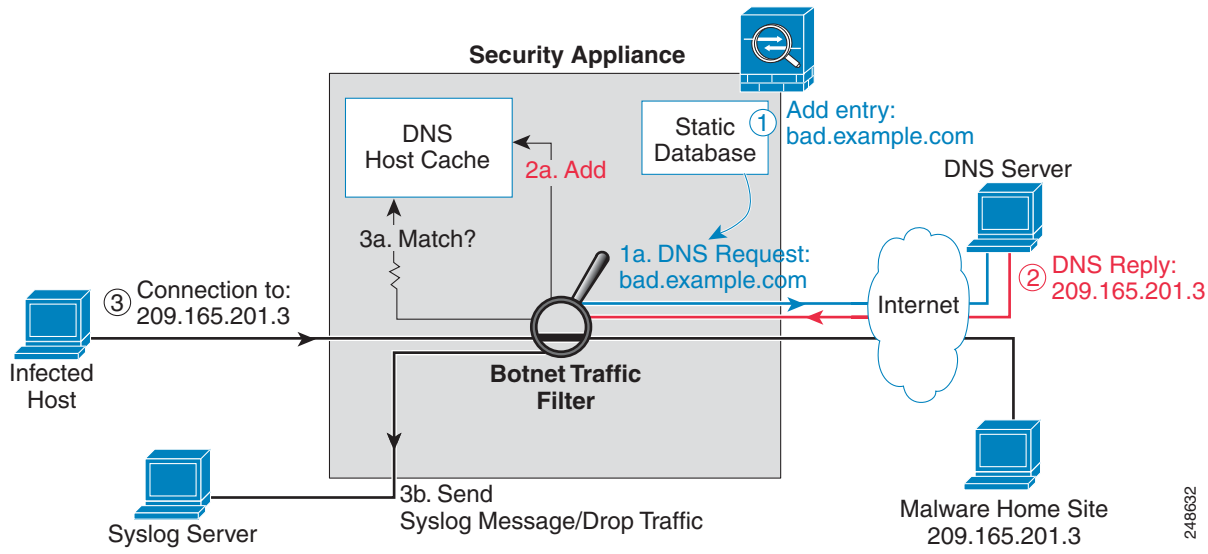

Licensing Requirements for the Botnet Traffic Filter

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	You need the following licenses: • Botnet Traffic Filter License. • Strong Encryption (3DES/AES) License to download the dynamic database.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support replication of the DNS reverse lookup cache, DNS host cache, or the dynamic database in Stateful Failover.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines and Limitations

- TCP DNS traffic is not supported.
- You can add up to 1000 blacklist entries and 1000 whitelist entries in the static database.
- The packet tracer is not supported.

Default Settings

By default, the Botnet Traffic Filter is disabled, as is use of the dynamic database.

For DNS inspection, which is enabled by default, Botnet Traffic Filter snooping is disabled by default.

Configuring the Botnet Traffic Filter

This section includes the following topics:

- Task Flow for Configuring the Botnet Traffic Filter, page 76-7

- Configuring the Dynamic Database, page 76-7
- Enabling DNS Snooping, page 76-9
- Adding Entries to the Static Database, page 76-9
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 76-10
- Blocking Botnet Traffic Manually, page 76-12
- Searching the Dynamic Database, page 76-13

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 76-7.

This procedure enables database updates from the Cisco update server, and also enables use of the downloaded dynamic database by the ASA. Disallowing use of the downloaded database is useful in multiple context mode so you can configure use of the database on a per-context basis.

Step 2 (Optional) Add static entries to the database. See the "Adding Entries to the Static Database" section on page 76-9.

This procedure lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. You might want to use the static database instead of the dynamic database if you do not want to download the dynamic database over the Internet.

Step 3 Enable DNS snooping. See the "Enabling DNS Snooping" section on page 76-9.

This procedure enables inspection of DNS packets, compares the domain name with those in the dynamic database or the static database (when a DNS server for the ASA is unavailable), and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

Step 4 Enable traffic classification and actions for the Botnet Traffic Filter. See the "Enabling Traffic Classification and Actions for the Botnet Traffic Filter" section on page 76-10.

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 76-12.

If you choose not to block malware traffic automatically, you can block traffic manually by configuring an access rule to deny traffic, or by using the **shun** command in the Command Line Interface tool to block all traffic to and from a host.

Configuring the Dynamic Database

This procedure enables database updates, and also enables use of the downloaded dynamic database by the ASA. 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 ASA use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, enable DNS per context.

Detailed Steps

Step 1 Enable downloading of the dynamic database.

- In Single mode, choose the **Configuration > Firewall > Botnet Traffic Filter > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.
- In multiple context mode in the System execution space, choose the **Configuration > Device Management > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.

This setting enables downloading of the dynamic database from the Cisco update server. In multiple context mode, enter this command in the system execution space. If you do not have a database already installed on the ASA, it downloads the database after approximately 2 minutes. The update server determines how often the ASA polls the server for future updates, typically every hour.

Step 2 (Multiple context mode only) In multiple context mode, click **Apply**. Then change to the context where you want to configure the Botnet Traffic Filter by double-clicking the context name in the Device List.

Step 3 In the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area, check the **Use Botnet data dynamically downloaded from updater server** check box.

Step 4 Click **Apply**.

Step 5 (Optional) If you want to later remove the database from running memory, perform the following steps:

a. Disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box.

b. Click **Apply**.

c. Click **Purge Botnet Database**.

d. To redownload the database, re-check the **Use Botnet data dynamically downloaded from updater server** check box.

e. Click **Apply**.

Note The Fetch Botnet Database button is for testing purposes only; it downloads and verifies the dynamic database, but does not store it in running memory.

For information about the Search Dynamic Database area, see the "Searching the Dynamic Database" section on page 76-13.

What to Do Next

See the "Adding Entries to the Static Database" section on page 76-9.

Adding Entries to the Static Database

The static database lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. Static blacklist entries are always designated with a Very High threat level. See the "Information About the Static Database" section on page 76-3 for more information.

Prerequisites

- In multiple context mode, perform this procedure in the context execution space.
- Enable ASA use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, enable DNS per context.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Botnet Traffic Filter > Black** or White List pane, click **Add** for the Whitelist or Blacklist.

The Enter hostname or IP Address dialog box appears.

Step 2 In the Addresses field, enter one or more domain names, IP addresses, and IP address/netmasks.

Enter multiple entries separated by commas, spaces, lines, or semi-colons. You can enter up to 1000 entries for each type.

Step 3 Click **OK**.

Step 4 Click **Apply**.

What to Do Next

See the "Enabling DNS Snooping" section on page 76-9.

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 58-1 and Chapter 44, "Configuring a Service Policy," for detailed information about configuring advanced DNS inspection options using the Modular Policy Framework.

Note You can also configure DNS snooping directly in the Configuration > Firewall > Service Policy Rules > Rule Actions > Protocol Inspection > Select DNS Inspect Map dialog box by checking the **Enable Botnet traffic filter DNS snooping** check box.

Restrictions

TCP DNS traffic is not supported.

Default DNS Inspection Configuration and Recommended Configuration

The default configuration for DNS inspection inspects all UDP DNS traffic on all interfaces, and does not have DNS snooping enabled.

We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the ASA.

For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Botnet Traffic Filter > DNS Snooping** pane.

All existing service rules that include DNS inspection are listed in the table.

Step 2 For each rule for which you want to enable DNS snooping, in the DNS Snooping Enabled column, check the check box.

Step 3 Click **Apply**.

What to Do Next

See the "Enabling Traffic Classification and Actions for the Botnet Traffic Filter" section on page 76-10.

Enabling Traffic Classification and Actions for the Botnet Traffic Filter

This procedure enables the Botnet Traffic Filter. The Botnet Traffic Filter compares the source and destination IP address in each initial connection packet to the following:

- Dynamic database IP addresses
- Static database IP addresses
- DNS reverse lookup cache (for dynamic database domain names)
- DNS host cache (for static database domain names)

When an address matches, the ASA sends a syslog message. The only additional action currently available is to drop the connection.

Prerequisites

In multiple context mode, perform this procedure in the context execution space.

Recommended Configuration

Although DNS snooping is not required, we recommend configuring DNS snooping for maximum use of the Botnet Traffic Filter (see the "Enabling DNS Snooping" section on page 76-9). 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 ASA.

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 "Adding ACLs and ACEs" section on page 26-2 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 Types" section on page 76-2 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 76-12.

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 ASA.

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 "Adding ACLs and ACEs" section on page 26-2 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 76-10), 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 96, "Configuring Logging," and Chapter 51, "Configuring Access Rules," for more information about creating an access rule.

Note If you create a reverse access rule form a Botnet Traffic Filter syslog message, and you do not have any other access rules applied to the interface, then you might inadvertently block all traffic. Normally, without an access rule, all traffic from a high security to a low security interface is allowed. But when you apply an access rule, all traffic is denied except traffic that you explicitly permit. Because the reverse access rule is a deny rule, be sure to edit the resulting access policy for the interface to permit other traffic.

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 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 ASA. This section includes the following topics:

- Botnet Traffic Filter Syslog Messaging, page 76-14
- Botnet Traffic Filter Monitor Panes, page 76-15

Botnet Traffic Filter Syslog Messaging

The Botnet Traffic Filter generates detailed syslog messages numbered 338*nnn*. Messages differentiate between incoming and outgoing connections, blacklist, whitelist, or greylist addresses, and many other variables. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.)

See the syslog messages guide for detailed information about syslog messages.

For the following syslog messages, a reverse access rule can be automatically created from the Real Time Log Viewer:

- 338001, 338002, 338003, 338004 (blacklist)
- 338201, 338202 (greylist)

See Chapter 96, “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 classified as whitelist, blacklist, and greylist connections, and how many connections were dropped. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The Details button shows how many packets at each threat level were classified or dropped.
Monitoring > Botnet Traffic Filter > Real-time Reports	Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected. If you right-click a site IP address, you can invoke the whois tool to learn more about the malware site. Reports can be saved as a PDF file.
Monitoring > Botnet Traffic Filter > Infected Hosts	Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The Maximum Connections option shows the 20 infected hosts with the most number of connections. The Latest Activity option shows the 20 hosts with the most recent activity. The Highest Threat Level option shows the 20 hosts that connected to the malware sites with the highest threat level. The Subnet option shows up to 20 hosts within the specified subnet. Reports can be saved as a PDF file, as either the Current View or the Whole Buffer. The Whole Buffer option shows all buffered infected-hosts information.
Monitoring > Botnet Traffic Filter > Updater Client	Shows information about the updater server, including the server IP address, the next time the ASA will connect with the server, and the database version last installed.
Monitoring > Botnet Traffic Filter > DNS Snooping	Shows the Botnet Traffic Filter DNS snooping actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.

Command	Purpose
Monitoring > Botnet Traffic Filter > Dynamic Database	Shows information about the dynamic database, including when the dynamic database was last downloaded, the version of the database, how many entries the database contains, and 10 sample entries.
Monitoring > Botnet Traffic Filter > ASP Table Hits	Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

Where to Go Next

- To configure the syslog server, see Chapter 96, "Configuring Logging."
- To block connections with an access rule, see Chapter 51, "Configuring Access Rules."

Feature History for the Botnet Traffic Filter

Table 76-2 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 76-2 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 1 hour timeout for reports for top hosts was removed; there is now no timeout. The following screens were introduced or modified: Configuration > Firewall > Botnet Traffic Filter > Traffic Settings, and Monitoring > Botnet Traffic Filter > Infected Hosts.

Configuring Threat Detection

This chapter describes how to configure threat detection statistics and scanning threat detection and includes the following sections:

- Information About Threat Detection, page 77-1
- Licensing Requirements for Threat Detection, page 77-1
- Configuring Basic Threat Detection Statistics, page 77-2
- Configuring Advanced Threat Detection Statistics, page 77-5
- Configuring Scanning Threat Detection, page 77-8

Information About Threat Detection

The threat detection feature consists of the following elements:

- Different levels of statistics gathering for various threats.

  Threat detection statistics can help you manage threats to your ASA; for example, if you enable scanning threat detection, then viewing statistics can help you analyze the threat. You can configure two types of threat detection statistics:

  - Basic threat detection statistics—Includes information about attack activity for the system as a whole. Basic threat detection statistics are enabled by default and have no performance impact.
  - Advanced threat detection statistics—Tracks activity at an object level, so the ASA can report activity for individual hosts, ports, protocols, or access lists. Advanced threat detection statistics can have a major performance impact, depending on the statistics gathered, so only the access list statistics are enabled by default.

- Scanning threat detection, which determines when a host is performing a scan.

  You can optionally shun any hosts determined to be a scanning threat.

Licensing Requirements for Threat Detection

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Configuring Basic Threat Detection Statistics

Basic threat detection statistics include activity that might be related to an attack, such as a DoS attack.

This section includes the following topics:

- Information About Basic Threat Detection Statistics, page 77-2
- Guidelines and Limitations, page 77-3
- Default Settings, page 77-3
- Configuring Basic Threat Detection Statistics, page 77-4
- Monitoring Basic Threat Detection Statistics, page 77-4
- Feature History for Basic Threat Detection Statistics, page 77-5

Information About Basic Threat Detection Statistics

Using basic threat detection statistics, the ASA monitors the rate of dropped packets and security events due to the following reasons:

- Denial by 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 77-8) takes this scanning attack rate information and acts on it by classifying hosts as attackers and automatically shunning them, for example.)
- Incomplete session detection such as TCP SYN attack detected or no data UDP session attack detected

When the ASA detects a threat, it immediately sends a system log message (733100). The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst rate interval is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each received event, the ASA checks the average and burst rate limits; if both rates are exceeded, then the ASA sends two separate system messages, with a maximum of one message for each rate type per burst period.

Basic threat detection affects performance only when there are drops or potential threats; even in this scenario, the performance impact is insignificant.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines

Supported in single mode only. Multiple mode is not supported.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Types of Traffic Monitored

Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.

Default Settings

Basic threat detection statistics are enabled by default.

Table 77-1 lists the default settings. You can view all these default settings using the **show running-config all threat-detection** command in Tools > Command Line Interface.

Table 77-1 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 20 second period.
	80 drops/sec over the last 3600 seconds.	320 drops/sec over the last 120 second period.
Scanning attack detected	5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
	4 drops/sec over the last 3600 seconds.	8 drops/sec over the last 120 second period.
Incomplete session detected such as TCP SYN attack detected or no data UDP session attack detected (combined)	100 drops/sec over the last 600 seconds.	200 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	160 drops/sec over the last 120 second period.
Denial by access lists	400 drops/sec over the last 600 seconds.	800 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	640 drops/sec over the last 120 second period.
• Basic firewall checks failed • Packets failed application inspection	400 drops/sec over the last 600 seconds.	1600 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	1280 drops/sec over the last 120 second period.

Table 77-1 Basic Threat Detection Default Settings (continued)

	Trigger Settings	
Packet Drop Reason	**Average Rate**	**Burst Rate**
Interface overload	2000 drops/sec over the last 600 seconds.	8000 drops/sec over the last 20 second period.
	1600 drops/sec over the last 3600 seconds.	6400 drops/sec over the last 120 second period.

Configuring Basic Threat Detection Statistics

This section describes how to configure basic threat detection statistics, including enabling or disabling it and changing the default limits.

Detailed Steps

Step 1 To enable or disable basic threat detection, choose the **Configuration > Firewall > Threat Detection** pane, and check the **Enable Basic Threat Detection** check box.

Step 2 Click **Apply**.

Monitoring Basic Threat Detection Statistics

To monitor basic threat detection statistics, perform the following task:

Path	Purpose
Home > Firewall Dashboard > Traffic Overview	Displays basic threat detection statistics. For a description of each event type, see the "Information About Basic Threat Detection Statistics" section on page 77-2.

Feature History for Basic Threat Detection Statistics

Table 77-2 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 77-2 Feature History for Basic Threat Detection Statistics

Feature Name	Platform Releases	Feature Information
Basic threat detection statistics	8.0(2)	Basic threat detection statistics was introduced. The following screen was introduced: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Traffic Overview.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

Configuring Advanced Threat Detection Statistics

You can configure the ASA to collect extensive statistics. This section includes the following topics:

- Information About Advanced Threat Detection Statistics, page 77-5
- Guidelines and Limitations, page 77-5
- Default Settings, page 77-6
- Configuring Advanced Threat Detection Statistics, page 77-6
- Monitoring Advanced Threat Detection Statistics, page 77-7
- Feature History for Advanced Threat Detection Statistics, page 77-8

Information About Advanced Threat Detection Statistics

Advanced threat detection statistics show both allowed and dropped traffic rates for individual objects such as hosts, ports, protocols, or access lists.

Caution Enabling advanced statistics can affect the ASA performance, depending on the type of statistics enabled. Enabling host statistics affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Port statistics, however, has modest impact.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines

Only TCP Intercept statistics are available in multiple mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Types of Traffic Monitored

Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.

Default Settings

By default, statistics for access lists are enabled.

Configuring Advanced Threat Detection Statistics

By default, statistics for access lists are enabled. To enable other statistics, perform the following steps.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Threat Detection** pane.

Step 2 In the Scanning Threat Statistics area, choose one of the following options:

- Enable *all* statistics—Click the **Enable All Statistics** radio button.
- Disable *all* statistics—Click the **Disable All Statistics** radio button.
- Enable only certain statistics—Click the **Enable Only Following Statistics** radio button.

Step 3 If you chose to **Enable Only Following Statistics**, then check one or more of the following check boxes:

- **Hosts**—Enables host statistics. The host statistics accumulate for as long as the host is active and in the scanning threat host database. The host is deleted from the database (and the statistics cleared) after 10 minutes of inactivity.
- **Access Rules** (enabled by default)—Enables statistics for access rules.
- **Port**—Enables statistics for TCP and UDP ports.
- **Protocol**—Enables statistics for non-TCP/UDP IP protocols.
- **TCP-Intercept**—Enables statistics for attacks intercepted by TCP Intercept (see the "Configuring Connection Settings" section on page 71-8 to enable TCP Intercept).

Step 4 For host, port, and protocol statistics, you can change the number of rate intervals collected. In the Rate Intervals area, choose **1 hour**, **1 and 8 hours**, or **1, 8 and 24 hours** for each statistics type. The default interval is **1 hour**, which keeps the memory usage low.

Step 5 For TCP Intercept statistics, you can set the following options in the TCP Intercept Threat Detection area:

- **Monitoring Window Size**—Sets the size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

- **Burst Threshold Rate**—Sets the threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, syslog message 733104 is generated.
- **Average Threshold Rate**—Sets the average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, syslog message 733105 is generated.

Click **Set Default** to restore the default values.

Step 6 Click **Apply**.

Monitoring Advanced Threat Detection Statistics

To monitor advanced threat detection statistics, perform one of the following tasks:

Path	Purpose
Home > Firewall Dashboard > Top 10 Access Rules Home > Firewall Dashboard > Top Usage Statistics	Displays the top 10 statistics. For the Top 10 Access Rules, permitted and denied traffic are not differentiated in this display. In the Traffic Overview > Dropped Packets Rate graph, you can track access list denies. The Top 10 Sources and Top 10 Destinations tabs show statistics for hosts. **Note**: Due to the threat detction algorithm, an interface used as a combination failover and state link could appear in the top 10 hosts; this is expected behavior, and you can ignore this IP address in the display. The Top 10 Services tab shows statistics for both ports and protocols (both must be enabled for the display), and shows the combined statistics of TCP/UDP port and IP protocol types. TCP (protocol 6) and UDP (protocol 17) are not included in the display for IP protocols; TCP and UDP ports are, however, included in the display for ports. If you only enable statistics for one of these types, port or protocol, then you will only view the enabled statistics. The Top Ten Protected Servers under SYN Attack area shows the TCP Intercept statistics. The display includes the top 10 protected servers under attack. The **detail** button shows history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds. From the Interval drop-down list, choose **Last 1 hour**, **Last 8 hour**, or **Last 24 hour**.

Feature History for Advanced Threat Detection Statistics

Table 77-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 77-3 Feature History for Advanced Threat Detection Statistics

Feature Name	Platform Releases	Feature Information
Advanced threat detection statistics	8.0(2)	Advanced threat detection statistics was introduced. The following screens were introduced: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Top 10 Access Rules, Home > Firewall Dashboard > Top Usage Status, Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack.
TCP Intercept statistics	8.0(4)/8.1(2)	TCP Intercept statistics were introduced. The following screens were introduced or modified: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack.
Customize host statistics rate intervals	8.1(2)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: Configuration > Firewall > Threat Detection.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Customize port and protocol statistics rate intervals	8.3(1)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: Configuration > Firewall > Threat Detection.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

Configuring Scanning Threat Detection

This section includes the following topics:

- Information About Scanning Threat Detection, page 77-9
- Guidelines and Limitations, page 77-9
- Default Settings, page 77-10
- Configuring Scanning Threat Detection, page 77-10

- Feature History for Scanning Threat Detection, page 77-11

Information About Scanning Threat Detection

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

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

If the scanning threat rate is exceeded, then the ASA sends a syslog message (733101), and optionally shuns the attacker. The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst event rate is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each event detected that is considered to be part of a scanning attack, the ASA checks the average and burst rate limits. If either rate is exceeded for traffic sent from a host, then that host is considered to be an attacker. If either rate is exceeded for traffic received by a host, then that host is considered to be a target.

Caution The scanning threat detection feature can affect the ASA performance and memory significantly while it creates and gathers host- and subnet-based data structure and information.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines

Supported in single mode only. Multiple mode is not supported.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Types of Traffic Monitored

- Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.
- Traffic that is denied by an access list does not trigger scanning threat detection; only traffic that is allowed through the ASA and that creates a flow is affected by scanning threat detection.

Default Settings

Table 77-4 lists the default rate limits for scanning threat detection.

Table 77-4 Default Rate Limits for Scanning Threat Detection

Average Rate	Burst Rate
5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
5 drops/sec over the last 3600 seconds.	10 drops/sec over the last 120 second period.

The burst rate is calculated as the average rate every *N* seconds, where *N* is the burst rate interval. The burst rate interval is 1/30th of the rate interval or 10 seconds, whichever is larger.

Configuring Scanning Threat Detection

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Threat Detection** pane, and check the **Enable Scanning Threat Detection** check box.

Step 2 (Optional) To automatically terminate a host connection when the ASA identifies the host as an attacker, check the **Shun Hosts detected by scanning threat** check box.

Step 3 (Optional) To except host IP addresses from being shunned, enter an address in the Networks excluded from shun field.

You can enter multiple addresses or subnets separated by commas. To choose a network from the list of IP address objects, click the **...** button.

Step 4 (Optional) To set the duration of a shun for an attacking host, check the **Set Shun Duration** check box and enter a value between 10 and 2592000 seconds. The default length is 3600 seconds (1 hour). To restore the default value, click **Set Default**.

Feature History for Scanning Threat Detection

Table 77-5 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 77-5 Feature History for Scanning Threat Detection

Feature Name	Platform Releases	Feature Information
Scanning threat detection	8.0(2)	Scanning threat detection was introduced. The following screen was introduced: Configuration > Firewall > Threat Detection.
Shun duration	8.0(4)/8.1(2)	You can now set the shun duration, The following screen was modified: Configuration > Firewall > Threat Detection.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

Using Protection Tools

This chapter describes some of the many tools available to protect your network and includes the following sections:

- Preventing IP Spoofing, page 78-1
- Configuring the Fragment Size, page 78-2
- Configuring TCP Options, page 78-3
- Configuring IP Audit for Basic IPS Support, page 78-5

Preventing IP Spoofing

This section 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 ASA only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the ASA 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 ASA, the ASA routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the ASA 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 ASA 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 ASA drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the ASA 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.

Configuration > Firewall > Advanced > Anti-Spoofing 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.

Configuring the Fragment Size

By default, the ASA 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 ASA. Fragmented packets are often used as DoS attacks.

To modify the IP fragment database parameters of an interface, perform the following steps:

Step 1 Choose the **Configuration > Firewall > Advanced > Fragment** pane, choose the interface to change in the Fragment table, and click **Edit**.

The Edit Fragment dialog box appears.

Step 2 In the Size field, set the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200.

Step 3 In the Chain field, set the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.

Step 4 In the Timeout field, set 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.

Step 5 Click **OK**.

Step 6 Click **Apply**.

Step 7 To view the fragment statistics, click **Show Fragment**. See the "Show Fragment" section on page 78-2 for more information.

Show Fragment

The Configuration > Properties > Fragment > Show Fragment pane displays the current IP fragment database statistics for each interface.

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.

Configuring TCP Options

The Configuration > Properties > 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 ASA to send TCP resets for all inbound TCP sessions that attempt to transit the ASA and are denied by the ASA based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the ASA silently discards denied packets.
  - Outbound Reset—Shows the interface reset setting for outbound TCP traffic, Yes or No. Enabling this setting causes the ASA to send TCP resets for all outbound TCP sessions that attempt to transit the ASA and are denied by the ASA based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the ASA 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 ASA based on access lists or AAA settings. When this option is not enabled, the ASA 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-ASA traffic as well as through the ASA 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 ASA 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 ASA 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 ASA 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 ASA 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 ASA 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 ASA 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.

TCP Reset Settings

The Configuration > Properties > TCP Options > TCP Reset Settings 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 ASA and are denied by the ASA based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the ASA 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 ASA and are denied by the ASA based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the ASA 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.

Configuring IP Audit for Basic IPS Support

The IP audit feature provides basic IPS support for the ASA that does not have an AIP SSM. It supports a basic list of signatures, and you can configure the ASA to perform one or more actions on traffic that matches a signature.

This section includes the following topics:

- IP Audit Policy, page 78-5
- Add/Edit IP Audit Policy Configuration, page 78-5
- IP Audit Signatures, page 78-6
- IP Audit Signature List, page 78-6

IP Audit Policy

The Configuration > Firewall > Advanced > IP Audit > 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.

Add/Edit IP Audit Policy Configuration

The Configuration > Firewall > Advanced > IP Audit > IP Audit Policy > 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.

IP Audit Signatures

The Configuration > Firewall > Advanced > IP Audit > 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 the "IP Audit Signature List" section on page 78-6.

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.

IP Audit Signature List

Table 78-1 lists supported signatures and system message numbers.

Table 78-1 Signature IDs and System Message Numbers

Signature ID	Message Number	Signature Title	Signature Type	Description
1000	400000	IP options-Bad Option List	Informational	Triggers on receipt of an IP datagram where the list of IP options in the IP datagram header is incomplete or malformed. The IP options list contains one or more options that perform various network management or debugging tasks.
1001	400001	IP options-Record Packet Route	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 7 (Record Packet Route).

Table 78-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
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.
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).

Table 78-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
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).
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).

Table 78-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
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.
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.

Table 78-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
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.
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.

Table 78-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
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 Filtering Services

This chapter describes how to use filtering services to provide greater control over traffic passing through the ASA and includes the following sections:

- Information About Web Traffic Filtering, page 79-1
- Configuring Filtering Rules, page 79-6
- Filtering the Rule Table, page 79-11
- Defining Queries, page 79-12
- Filtering URLs and FTP Requests with an External Server, page 79-2

Information About Web Traffic Filtering

You can use web traffic filtering in two distinct ways:

- Filtering ActiveX objects or Java applets
- Filtering with an external filtering server

Instead of blocking access altogether, you can remove specific undesirable objects from web traffic, such as ActiveX objects or Java applets, that may pose a security threat in certain situations.

You can use web traffic filtering to direct specific traffic to an external filtering server, such an Secure Computing SmartFilter (formerly N2H2) or the Websense filtering server. You can enable long URL, HTTPS, and FTP filtering using either Websense or Secure Computing SmartFilter for web traffic filtering. Filtering servers can block traffic to specific sites or types of sites, as specified by the security policy.

Note URL caching will only work if the version of the URL server software from the URL server vendor supports it.

Because web traffic 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 web traffic filtering server, the time required for the initial connection may be noticeably slower when filtering traffic with an external filtering server.

Model	License Requirement
All models	Base License.

Filtering URLs and FTP Requests with an External Server

This section describes how to filter URLs and FTP requests with an external server and includes the following topics:

- Information About URL Filtering, page 79-2
- Licensing Requirements for URL Filtering, page 79-3
- Guidelines and Limitations for URL Filtering, page 79-3
- Identifying the Filtering Server, page 79-3
- Configuring Additional URL Filtering Settings, page 79-4
- Feature History for URL Filtering, page 79-12

Information About 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 ASA performance by using a separate server running one of the following Internet filtering products:

- Websense Enterprise for filtering HTTP, HTTPS, and FTP.
- McAfee SmartFilter (formerly N2H2) for filtering HTTP, HTTPS, FTP, and long URL filtering.

  In long URLs, the URL in the Referer field might contain a “host:” text string, which could cause the HTTP GET header to be incorrectly parsed as containing the HTTP Host parameter. The ASA, however, correctly parses the Referer field even when it contains a “host:” text string and forwards the header to the McAfee SmartFilter server with the correct Referer URL.

Note URL caching will only work if the version of the URL server software from the URL server vendor supports it.

Although ASA performance is less affected when using an external server, you might notice longer access times to websites or FTP servers when the filtering server is remote from the ASA.

When filtering is enabled and a request for content is directed through the ASA, 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 ASA forwards the response from the content server to the originating client. If the filtering server denies the connection, the ASA drops the response and sends a message or return code indicating that the connection was not successful.

If user authentication is enabled on the ASA, then the ASA also sends the username to the filtering server. The filtering server can use user-specific filtering settings or provide enhanced reporting about usage.

Licensing Requirements for URL Filtering

The following table shows the licensing requirements for URL filtering:

Table 79-1 Licensing Requirements

Model	License Requirement
All models	Base License.

Guidelines and Limitations for URL Filtering

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

Does not support IPv6.

Identifying the Filtering Server

You can identify up to four filtering servers per context. The ASA uses the servers in order until a server responds. In single mode, a maximum of 16 of the same type of filtering servers are allowed. 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 or HTTPS.

To specify the external filtering server, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration > Firewall > URL Filtering Servers**.

Step 2 In the URL Filtering Server Type area, click one of the following options:

- **Websense**
- **Secure Computing SmartFilter**

Step 3 If you chose the second option, enter the Secure Computing SmartFilter port number if it is different than the default port number, which is 4005.

Step 4 In the URL Filtering Servers area, click **Add**.

If you chose the Websense option, the Add Parameters for Websense URL Filtering dialog box appears.

- Choose the interface on which the URL filtering server is connected from the drop-down list.
- Enter the IP address of the URL filtering server.

- Enter the number of seconds after which the request to the URL filtering server times out. The default is 30 seconds.
- In the Protocol area, to specify which TCP version to use to communicate with the URL filtering server, click one of the following radio buttons:
  - TCP 1
  - TCP 4
  - UDP 4
- Enter the maximum number of TCP connections allowed for communicating with the URL filtering server, and click **OK**.

The new Websense URL filtering server properties appear in the URL Filtering Servers pane. To change these properties, click **Edit**. To add more Websense URL filtering servers after you have added the first Websense URL filtering server, click **Add** or **Insert**. To remove a Websense URL filtering server, click **Delete**.

If you chose the Secure Computing SmartFilter URL Filtering option, the Add Parameters for Secure Computing SmartFilter URL Filtering dialog box appears.

- Choose the interface on which the URL filtering server is connected from the drop-down list.
- Enter the IP address of the URL filtering server.
- Enter the number of seconds after which the request to the URL filtering server times out. The default is 30 seconds.
- In the Protocol area, to specify which protocol type to use to communicate with the URL filtering server, click one of the following radio buttons:
  - TCP
  - UDP
- Enter the maximum number of TCP connections allowed for communicating with the URL filtering server, and click **OK**.

The new Secure Computing SmartFilter URL filtering server properties appear in the URL Filtering Servers pane. To change these properties, click **Edit**. To add more Secure Computing SmartFilter URL filtering servers after you have defined the first Secure Computing SmartFilter URL filtering server, click **Add** or **Insert**. To remove a Secure Computing SmartFilter URL filtering server, click **Delete**.

Configuring Additional URL Filtering Settings

After you have accessed a website, the filtering server can allow the ASA to cache the server address for a certain period of time, as long as each website hosted at the address is in a category that is permitted at all times. When you access the server again, or if another user accesses the server, the ASA does not need to consult the filtering server again to obtain the server address.

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.

This section describes how to configure additional URL filtering settings and includes the following topics:

- Buffering the Content Server Response, page 79-5
- Caching Server Addresses, page 79-5
- Filtering HTTP URLs, page 79-6

Buffering the Content Server Response

When you issue a request to connect to a content server, the ASA 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 behavior delays the web server response for the web client, because the web 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 behavior prevents the delay that might otherwise occur.

To configure buffering for responses to HTTP or FTP requests, perform the following steps:

Step 1 In the URL Filtering Servers pane, click **Advanced** to display the Advanced URL Filtering dialog box.

Step 2 In the URL Buffer Size area, check the **Enable buffering** check box.

Step 3 Enter the number of 1550-byte buffers. Valid values range from 1 to 128.

Step 4 Click **OK** to close this dialog box.

Caching Server Addresses

After you access a website, the filtering server can allow the ASA to cache the server address for a certain period of time, as long as each website hosted at the address is in a category that is permitted at all times. When you access the server again, or if another user accesses the server, the ASA 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. You can accumulate Websense run logs before using the **url-cache** command.

To improve throughput, perform the following steps:

Step 1 In the URL Filtering Servers pane, click **Advanced** to display the Advanced URL Filtering dialog box.

Step 2 In the URL Cache Size area, check the **Enable caching based on** check box to enable caching according to the specified criteria.

Step 3 Click one of the following radio buttons:

- Destination Address—This option caches entries according to the URL destination address. Choose this setting if all users share the same URL filtering policy on the Websense server.
- Source/Destination Address—This option caches entries according to both the source address that initiates the URL request and the URL destination address. Choose this setting if users do not share the same URL filtering policy on the server.

Step 4 Enter the cache size within the range from 1 to 128 (KB).

Step 5 Click **OK** to close this dialog box.

Filtering HTTP URLs

This section describes how to configure HTTP filtering with an external filtering server and includes the following topics:

- Enabling Filtering of Long HTTP URLs, page 79-6

Enabling Filtering of Long HTTP URLs

By default, the ASA considers an HTTP URL to be a long URL if it is greater than 1159 characters. You can increase the maximum length allowed.

To configure the maximum size of a single URL, perform the following steps:

Step 1 In the URL Filtering Servers pane, click **Advanced** to display the Advanced URL Filtering dialog box.

Step 2 In the Long URL Support area, check the **Use Long URL** check box to enable long URLs for filtering servers.

Step 3 Enter the maximum URL length allowed, up to a maximum of 4 KB.

Step 4 Enter the memory allocated for long URLs in KB.

Step 5 Click **OK** to close this dialog box.

Configuring Filtering Rules

Before you can add an HTTP, HTTPS, or FTP filter rule, you must enable a URL filtering server. To enable a URL filtering server, choose **Configuration > Firewall > URL Filtering Servers**.

To configure filtering rules, perform the following steps:

Step 1 From the ASDM main window, choose **Configuration > Firewall > Filter Rules**.

Step 2 In the toolbar, click **Add** to display the types of filter rules that are available to add from the following list:

- Add Filter ActiveX Rule
- Add Filter Java Rule
- Add Filter HTTP Rule
- Add Filter HTTPS Rule
- Add Filter FTP Rule

Step 3 If you chose Add Filter ActiveX Rule, specify the following settings:

- Click one of the following radio buttons: **Filter ActiveX** or **Do not filter ActiveX**.
- Enter the source of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any source address.

  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Source dialog box. Choose a host or address from the drop-down list.
- Enter the destination of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any destination address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Destination dialog box. Choose a host or address from the drop-down list.
- Identify the service of the traffic to which the filtering action applies. To identify the service, enter one of the following:
  - tcp/*port*—The port number can range 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.

    >—Greater than. For example, >tcp/2000.

    - —Range. For example, tcp/2000-3000.
  - Enter a well-known service name, such as HTTP or FTP.
  - Click the ellipses to display the Browse Service dialog box. Choose a service from the drop-down list.
- Click **OK** to close this dialog box.
- Click **Apply** to save your changes.

Step 4 If you chose Add Filter Java Rule, specify the following settings:

- Click one of the following radio buttons: **Filter Java** or **Do not filter Java**.
- Enter the source of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any source address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Source dialog box. Choose a host or address from the drop-down list.
- Enter the destination of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any destination address.
  - Enter a hostname.

  - 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**.
  - Click the ellipses to display the Browse Destination dialog box. Choose a host or address from the drop-down list.
- Identify the service of the traffic to which the filtering action applies. To identify the service, enter one of the following:
  - 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.

    >—Greater than. For example, >tcp/2000.

    - —Range. For example, tcp/2000-3000.
  - Enter a well-known service name, such as HTTP or FTP.
  - Click the ellipses to display the Browse Service dialog box. Choose a service from the drop-down list.
- Click **OK** to close this dialog box.
- Click **Apply** to save your changes.

Step 5 If you chose Add Filter HTTP Rule, specify the following settings:

- Click one of the following radio buttons: **Filter HTTP** or **Do not filter HTTP**.
- Enter the source of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any source address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Source dialog box. Choose a host or address from the drop-down list.
- Enter the destination of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any destination address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Destination dialog box. Choose a host or address from the drop-down list.
- Identify the service of the traffic to which the filtering action applies. To identify the service, enter one of the following:
  - tcp/*port*—The port number can range 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.

>—Greater than. For example, >tcp/2000.

- —Range. For example, tcp/2000-3000.

– Enter a well-known service name, such as HTTP or FTP.

– Click the ellipses to display the Browse Service dialog box. Choose a service from the drop-down list.

- Choose the action to take when the URL exceeds the specified size from the drop-down list.
- Check the **Allow outbound traffic if URL server is not available check box** to connect without URL filtering being performed. When this check box is unchecked, you cannot connect to Internet websites if the URL server is unavailable.
- Check the **Block users from connecting to an HTTP proxy server check box to** prevent HTTP requests made through a proxy server.
- Check the **Truncate CGI parameters from URL sent to URL server** check box to have the ASA forward only the CGI script location and the script name, without any parameters, to the filtering server.
- Click **OK** to close this dialog box.
- Click **Apply** to save your changes.

Step 6 If you chose Add Filter HTTPS Rule, specify the following settings:

- Click one of the following radio buttons: **Filter HTTPS** or **Do not filter HTTPS**.
- Enter the source of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  – Enter **any** to indicate any source address.
  – Enter a hostname.
  – 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**.
  – Click the ellipses to display the Browse Source dialog box. Choose a host or address from the drop-down list.
- Enter the destination of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  – Enter **any** to indicate any destination address.
  – Enter a hostname.
  – 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**.
  – Click the ellipses to display the Browse Destination dialog box. Choose a host or address from the drop-down list.
- Identify the service of the traffic to which the filtering action applies. To identify the service, enter one of the following:
  – tcp/*port*—The port number can range 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.

    >—Greater than. For example, >tcp/2000.

    - —Range. For example, tcp/2000-3000.

  - Enter a well-known service name, such as HTTP or FTP.
  - Click the ellipses to display the Browse Service dialog box. Choose a service from the drop-down list.
- Check the **Allow outbound traffic if URL server is not available check box** to connect without URL filtering being performed. When this check box is unchecked, you cannot connect to Internet websites if the URL server is unavailable.
- Click **OK** to close this dialog box.
- Click **Apply** to save your changes.

Step 7 If you chose Add Filter FTP Rule, specify the following settings:

- Click one of the following radio buttons: **Filter FTP** or **Do not filter FTP**.
- Enter the source of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any source address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Source dialog box. Choose a host or address from the drop-down list.
- Enter the destination of the traffic to which the filtering action applies. To enter the source, choose from the following options:
  - Enter **any** to indicate any destination address.
  - Enter a hostname.
  - 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**.
  - Click the ellipses to display the Browse Destination dialog box. Choose a host or address from the drop-down list.
- Identify the service of the traffic to which the filtering action applies. To identify the service, enter one of the following:
  - tcp/*port*—The port number can range 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.

    >—Greater than. For example, >tcp/2000.

    - —Range. For example, tcp/2000-3000.
  - Enter a well-known service name, such as http or ftp.
  - Click the ellipses to display the Browse Service dialog box. Choose a service from the drop-down list.
- Check the **Allow outbound traffic if URL server is not available check box** to connect without URL filtering being performed. When this check box is unchecked, you cannot connect to Internet websites if the URL server is unavailable.
- Check the **Block interactive FTP sessions (block if absolute FTP path is not provided)** check box to drop FTP requests if they use a relative path name to the FTP directory.

- Click **OK** to close this dialog box.
- Click **Apply** to save your changes.

Step 8 To modify a filtering rule, select it and click **Edit** to display the Edit Filter Rule dialog box for the specified filtering rule.

Step 9 Make the required changes, then click **OK** to close this dialog box.

Step 10 Click **Apply** to save your changes.

Filtering the Rule Table

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 to display the Filter toolbar.

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 composed of source, destination, service, and rule type information.

Step 3 For Source, Destination, Source or Destination, and Service filters, perform the following steps:

a. Enter the string to match using one of the following methods:

  – Type the source, destination, or service name in the adjacent field.

  – Click the ellipses to open a Browse dialog box from which you can choose existing services, IP addresses, or host names.

b. Choose the match criteria from the drop-down list. Choose **is** for exact string matches or **contains** for partial string matches.

Step 4 For Rule Type filters, choose the rule type from the list.

Step 5 For Query filters, click **Define Query**. To define queries, see the "Defining Queries" section on page 79-12.

Step 6 To apply the filter to the rule table, click **Filter**.

Step 7 To remove the filter from the rule table and display all rule entries, click **Clear**.

Step 8 To show the packet trace for the selected rule, click **Packet Trace**.

Step 9 To show and hide the selected rule diagram, click **Diagram**.

Step 10 To remove a filter rule and place it elsewhere, click **Cut**.

Step 11 To copy a filter rule, click **Copy**. Then to move the copied filter rule elsewhere, click **Paste**.

Step 12 To delete a selected filter rule, click **Delete**.

Defining Queries

To define queries, perform the following steps:

Step 1 Enter the IP address or hostname of the source. Choose **is** for an exact match or choose **contains** for a partial match. Click the ellipses to display the Browse Source dialog box. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them with commas.

Step 2 Enter the IP address or hostname of the destination. Choose **is** for an exact match or choose **contains** for a partial match. Click the ellipses to display the Browse Destination dialog box. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them with commas.

Step 3 Enter the IP address or hostname of the source or destination. Choose **is** for an exact match or choose **contains** for a partial match. Click the ellipses to display the Browse Source dialog box. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them with commas.

Step 4 Enter the protocol, port, or name of a service. Choose **is** for an exact match or choose **contains** for a partial match. Click the ellipses to display the Browse Service dialog box. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them with commas.

Step 5 Choose the rule type from the drop-down list.

Step 6 Click **OK** to close this dialog box.

After you click **OK**, the filter is immediately applied to the rule table. To remove the filter, click **Clear**.

Feature History for URL Filtering

Table 79-2 lists the release history for URL filtering. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 79-2 Feature History for URL Filtering

Feature Name	Platform Releases	Feature Information
URL filtering	7.0(1)	Filters URLs based on an established set of filtering criteria.

PART 14

Configuring Modules

Configuring the ASA IPS Module

This chapter describes how to configure the ASA IPS module. The ASA IPS module might be a physical module or a software module, depending on your ASA model. For a list of supported ASA IPS modules per ASA model, see the *Cisco ASA Compatibility Matrix*:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

This chapter includes the following sections:

- Information About the ASA IPS module, page 81-1
- Licensing Requirements for the ASA IPS module, page 81-5
- Guidelines and Limitations, page 81-5
- Default Settings, page 81-6
- Configuring the ASA IPS module, page 81-7
- Managing the ASA IPS module, page 81-19
- Monitoring the ASA IPS module, page 81-24
- Feature History for the ASA IPS module, page 81-24

Information About the ASA IPS module

The ASA IPS module runs advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network. This section includes the following topics:

- How the ASA IPS module Works with the ASA, page 81-2
- Operating Modes, page 81-3
- Using Virtual Sensors (ASA 5510 and Higher), page 81-3
- Information About Management Access, page 81-4

How the ASA IPS module Works with the ASA

The ASA IPS module runs a separate application from the ASA. The ASA IPS module might include an external management interface so you can connect to the ASA IPS module directly; if it does not have a management interface, you can connect to the ASA IPS module through the ASA interface. The ASA IPS SSP on the ASA 5585-X includes data interfaces; these interfaces provide additional port-density for the ASA. However, the overall through-put of the ASA is not increased.

Traffic goes through the firewall checks before being forwarded to the ASA IPS module. When you identify traffic for IPS inspection on the ASA, traffic flows through the ASA and the ASA IPS module as follows. **Note**: This example is for "inline mode." See the "Operating Modes" section on page 81-3 for information about "promiscuous mode," where the ASA only sends a copy of the traffic to the ASA IPS module.

1. Traffic enters the ASA.
2. Incoming VPN traffic is decrypted.
3. Firewall policies are applied.
4. Traffic is sent to the ASA IPS module.
5. The ASA IPS module applies its security policy to the traffic, and takes appropriate actions.
6. Valid traffic is sent back to the ASA; the ASA IPS module might block some traffic according to its security policy, and that traffic is not passed on.
7. Outgoing VPN traffic is encrypted.
8. Traffic exits the ASA.

Figure 81-1 shows the traffic flow when running the ASA IPS module in inline mode. In this example, the ASA IPS module automatically blocks traffic that it identified as an attack. All other traffic is forwarded through the ASA.

Figure 81-1 ASA IPS module Traffic Flow in the ASA: Inline Mode

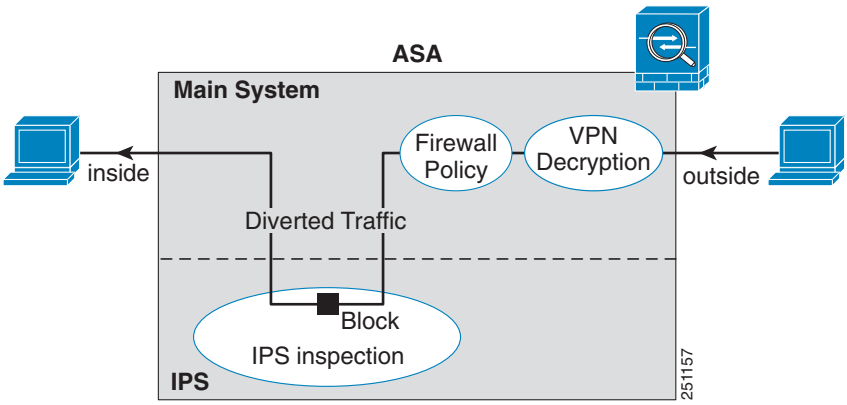

Operating Modes

You can send traffic to the ASA IPS module using one of the following modes:

- Inline mode—This mode places the ASA IPS module directly in the traffic flow (see Figure 81-1). No traffic that you identified for IPS inspection can continue through the ASA without first passing through, and being inspected by, the ASA IPS module. This mode is the most secure because every packet that you identify for inspection is analyzed before being allowed through. Also, the ASA IPS module can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput.
- Promiscuous mode—This mode sends a duplicate stream of traffic to the ASA IPS module. This mode is less secure, but has little impact on traffic throughput. Unlike inline mode, in promiscuous mode the ASA IPS module can only block traffic by instructing the ASA to shun the traffic or by resetting a connection on the ASA. Also, while the ASA IPS module is analyzing the traffic, a small amount of traffic might pass through the ASA before the ASA IPS module can shun it. Figure 81-2 shows the ASA IPS module in promiscuous mode. In this example, the ASA IPS module sends a shun message to the ASA for traffic it identified as a threat.

Figure 81-2 ASA IPS module Traffic Flow in the ASA: Promiscuous Mode

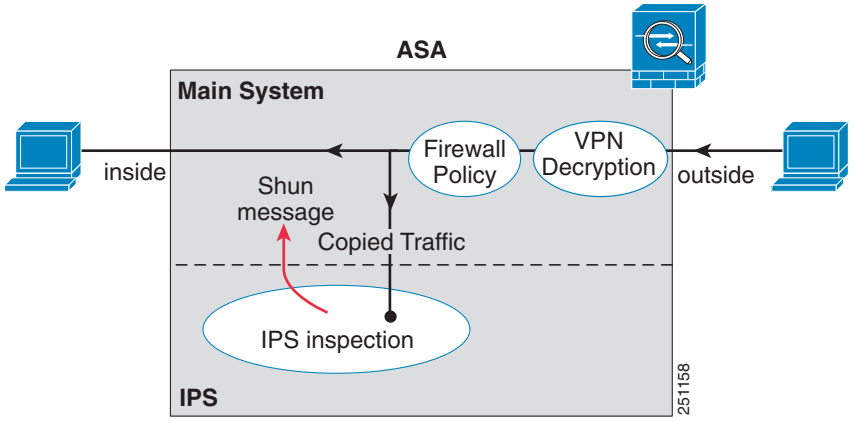

Using Virtual Sensors (ASA 5510 and Higher)

The ASA IPS module running IPS software Version 6.0 and later can run multiple virtual sensors, which means you can configure multiple security policies on the ASA IPS module. You can assign each ASA security context or single mode ASA to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. See the IPS documentation for more information about virtual sensors, including the maximum number of sensors supported.

Figure 81-3 shows one security context paired with one virtual sensor (in inline mode), while two security contexts share the same virtual sensor.

Figure 81-3 Security Contexts and Virtual Sensors

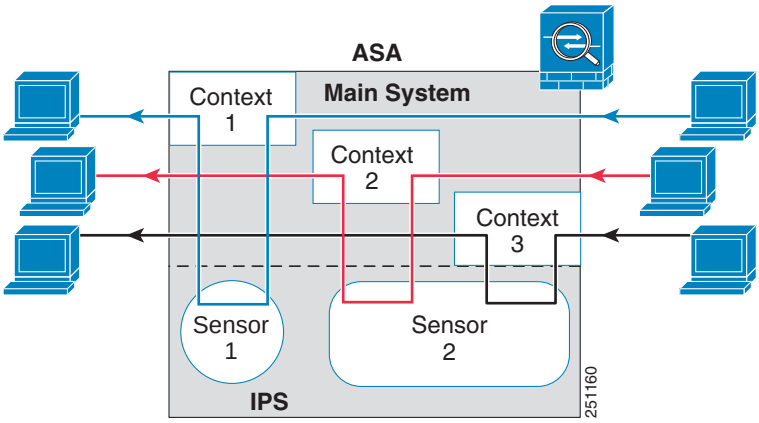

Figure 81-4 shows a single mode ASA paired with multiple virtual sensors (in inline mode); each defined traffic flow goes to a different sensor.

Figure 81-4 Single Mode ASA with Multiple Virtual Sensors

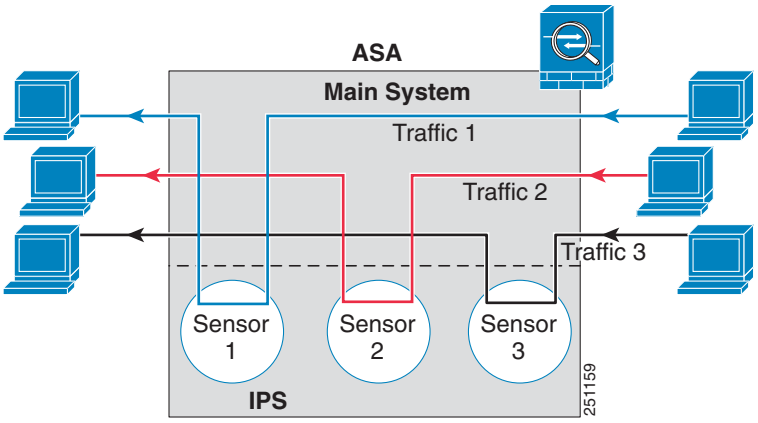

Information About Management Access

You can manage the IPS application using the following methods:

- Sessioning to the module from the ASA—If you have CLI access to the ASA, then you can session to the module and access the module CLI. See the "Sessioning to the Module from the ASA (May Be Required)" section on page 81-11.
- Connecting to the IPS management interface using ASDM or SSH—After you launch ASDM from the ASA, your management station connects to the module management interface to configure the IPS application. For SSH, you can access the module CLI directly on the module management interface. (Telnet access requires additional configuration in the module application). The module management interface can also be used for sending syslog messages or allowing updates for the module application, such as signature database updates.

See the following information about the management interface:

- ASA 5510, ASA 5520, ASA 5540, ASA 5580, ASA 5585-X—The IPS management interface is a separate external Gigabit Ethernet interface.
- ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X—These models run the ASA IPS module as a software module. The IPS management interface shares the Management 0/0 interface with the ASA. Separate MAC addresses and IP addresses are supported for the ASA and ASA IPS module. You must perform configuration of the IPS IP address within the IPS operating system (using the CLI or ASDM). However, physical characteristics (such as enabling the interface) are configured on the ASA. You can remove the ASA interface configuration (specifically the interface name) to dedicate this interface as an IPS-only interface. This interface is management-only.
- ASA 5505—You can use an ASA VLAN to allow access to an internal management IP address over the backplane.

Licensing Requirements for the ASA IPS module

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X	IPS Module License. **Note** The IPS module license lets you run the IPS software module on the ASA. You must also purchase a separate IPS signature subscription; for failover, purchase a subscription for each unit. To obtain IPS signature support, you must purchase the ASA with IPS pre-installed (the part number must include "IPS"). The combined failover cluster license does not let you pair non-IPS and IPS units. For example, if you buy the IPS version of the ASA 5515-X (part number ASA5515-IPS-K9) and try to make a failover pair with a non-IPS version (part number ASA5515-K9), then you will not be able to obtain IPS signature updates for the ASA5515-K9 unit, even though it has an IPS module license inherited from the other unit.
All other models	Base License.

The ASA IPS module 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 does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Model Guidelines

- See the *Cisco ASA Compatibility Matrix* for information about which models support which modules:

  http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

- The ASA 5505 does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.
- The ASA IPS module for the ASA 5510 and higher supports higher performance requirements, while the ASA IPS module for the ASA 5505 is designed for a small office installation. The following features are not supported for the ASA 5505:
  - Virtual sensors
  - Anomaly detection
  - Unretirement of default retired signatures

Additional Guidelines

- The total throughput for the ASA plus the IPS module is lower than ASA throughput alone.
  - ASA 5512-X through ASA 5555-X—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/qa_c67-700608.html
  - ASA 5585-X—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/qa_c67-617018.html
  - ASA 5505 through ASA 5540—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/product_data_sheet0900aecd802930c5.html
- You cannot change the software type installed on the module; if you purchase an ASA IPS module, you cannot later install other software on it.

Default Settings

Table 81-1 lists the default settings for the ASA IPS module.

Table 81-1 Default Network Parameters

Parameters	Default
Management VLAN (ASA 5505 only)	VLAN 1
Management IP address	192.168.1.2/24
Gateway	192.168.1.1/24 (the default ASA management IP address)
Username	cisco
Password	cisco

Note The default management IP address on the ASA is 192.168.1.1/24.

Configuring the ASA IPS module

This section describes how to configure the ASA IPS module and includes the following topics:

- Task Flow for the ASA IPS Module, page 81-7
- Connecting the ASA IPS Management Interface, page 81-8
- Configuring Basic IPS Module Network Settings, page 81-12
- (ASA 5512-X through ASA 5555-X) Booting the Software Module, page 81-12
- Configuring the Security Policy on the ASA IPS Module, page 81-15
- Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher), page 81-17
- Diverting Traffic to the ASA IPS module, page 81-18

Task Flow for the ASA IPS Module

Configuring the ASA IPS module is a process that includes configuration of the IPS security policy on the ASA IPS module and then configuration of the ASA to send traffic to the ASA IPS module. To configure the ASA IPS module, perform the following steps:

Step 1 Cable the ASA IPS management interface. See the "Connecting the ASA IPS Management Interface" section on page 81-8.

Step 2 Session to the module. Access the IPS CLI over the backplane. For ASDM users, you may need to session to the module to boot the IPS software if it is not running. See the "Sessioning to the Module from the ASA (May Be Required)" section on page 81-11.

Step 3 (ASA 5512-X through ASA 5555-X; may be required) Install the software module. See the "(ASA 5512-X through ASA 5555-X) Booting the Software Module" section on page 81-12.

Step 4 Depending on your ASA model:

- (ASA 5510 and higher) Configure basic network settings for the IPS module. See the "(ASA 5510 and Higher) Configuring Basic Network Settings" section on page 81-13.
- (ASA 5505) Configure the management VLAN and IP address for the IPS module. See the "(ASA 5505) Configuring Basic Network Settings" section on page 81-14.

Step 5 On the module, configure the inspection and protection policy, which determines how to inspect traffic and what to do when an intrusion is detected. See the "Configuring the Security Policy on the ASA IPS Module" section on page 81-15.

Step 6 (ASA 5510 and higher, optional) On the ASA in multiple context mode, specify which IPS virtual sensors are available for each context (if you configured virtual sensors). See the "Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher)" section on page 81-17.

Step 7 On the ASA, identify traffic to divert to the ASA IPS module. See the "Diverting Traffic to the ASA IPS module" section on page 81-18.

Connecting the ASA IPS Management Interface

In addition to providing management access to the IPS module, the IPS management interface needs access to an HTTP proxy server or a DNS server and the Internet so it can download global correlation, signature updates, and license requests. This section describes recommended network configurations. Your network may differ.

- ASA 5510, ASA 5520, ASA 5540, ASA 5580, ASA 5585-X (Physical Module), page 81-8
- ASA 5512-X through ASA 5555-X (Software Module), page 81-9
- ASA 5505, page 81-10

ASA 5510, ASA 5520, ASA 5540, ASA 5580, ASA 5585-X (Physical Module)

The IPS module includes a separate management interface from the ASA.

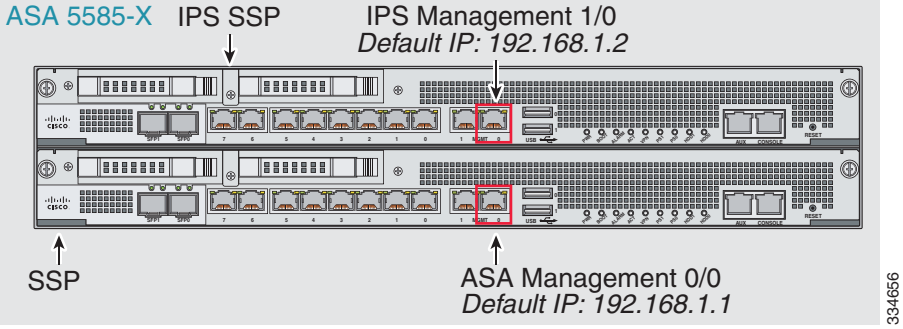

If you have an inside router

If you have an inside router, you can route between the management network, which can include both the ASA Management 0/0 and IPS Management 1/0 interfaces, and the ASA inside network. Be sure to also add a route on the ASA to reach the Management network through the inside router.

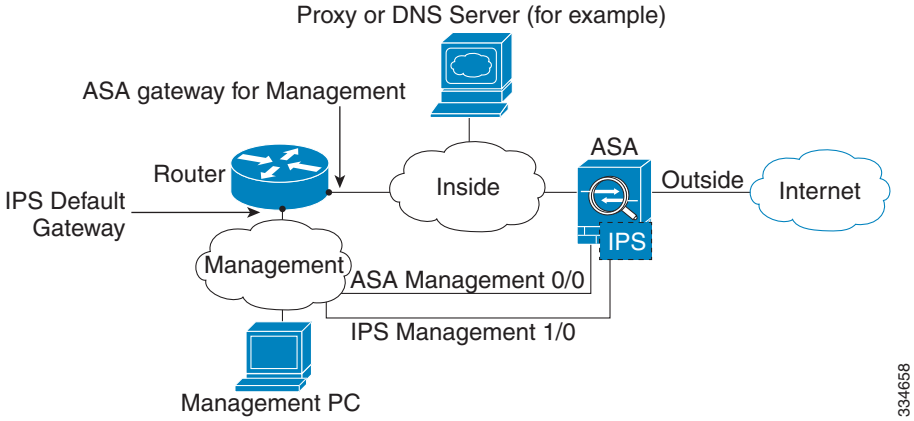

If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network, which would require an inside router to route between the networks. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. Because the IPS module is a separate device from the ASA, you can configure the IPS Management 1/0 address to be on the same network as the inside interface.

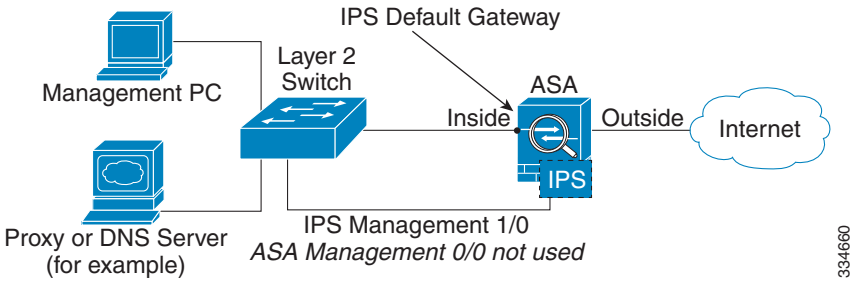

ASA 5512-X through ASA 5555-X (Software Module)

These models run the IPS module as a software module, and the IPS management interface shares the Management 0/0 interface with the ASA.

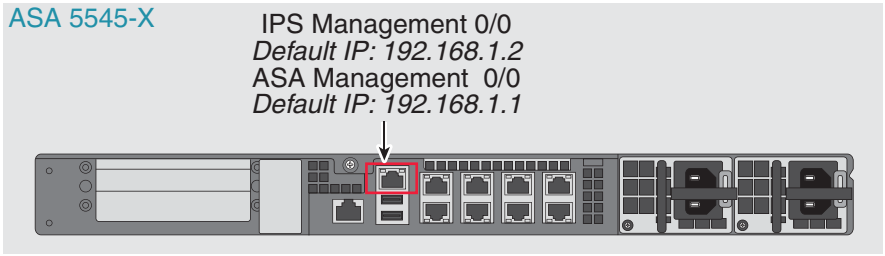

If you have an inside router

If you have an inside router, you can route between the Management 0/0 network, which includes both the ASA and IPS management IP addresses, and the inside network. Be sure to also add a route on the ASA to reach the Management network through the inside router.

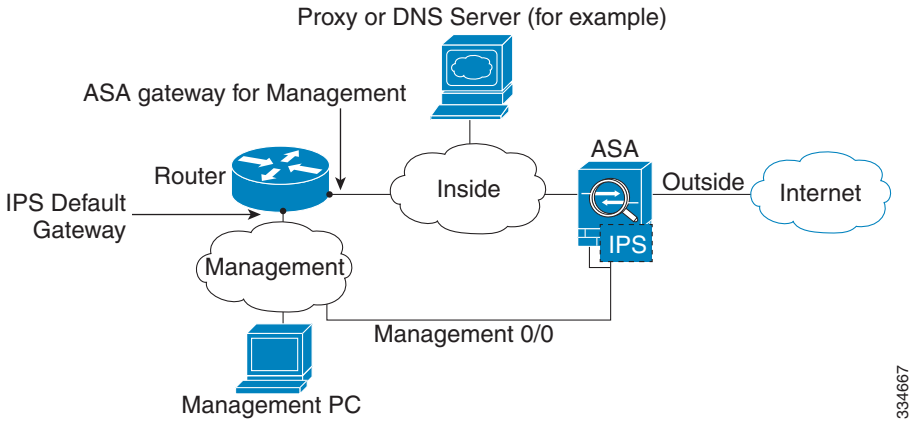

If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. If you remove the ASA-configured name from the Management 0/0 interface, you can still configure the IPS IP address for that interface. Because the IPS module is essentially a separate device from the ASA, you *can* configure the IPS management address to be on the same network as the inside interface.

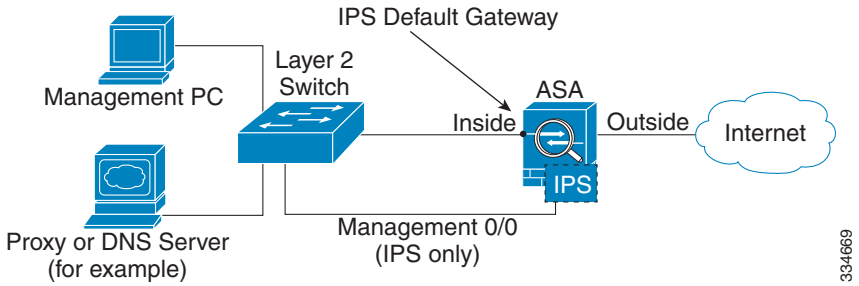

Note You must remove the ASA-configured name for Management 0/0; if it is configured on the ASA, then the IPS address must be on the same network as the ASA, and that excludes any networks already configured on other ASA interfaces. If the name is not configured, then the IPS address can be on any network, for example, the ASA inside network.

ASA 5505

The ASA 5505 does not have a dedicated management interface. You must use an ASA VLAN to access an internal management IP address over the backplane. Connect the management PC to one of the following ports: Ethernet 0/1 through 0/7, which are assigned to VLAN 1.

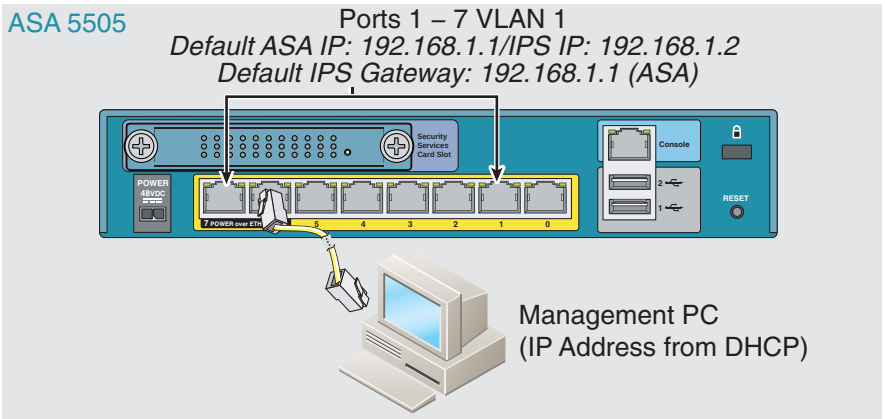

What to Do Next

- (ASA 5510 and higher) Configure basic network settings. See the "(ASA 5510 and Higher) Configuring Basic Network Settings" section on page 81-13.
- (ASA 5505) Configure management interface settings. See the "(ASA 5505) Configuring Basic Network Settings" section on page 81-14.

Sessioning to the Module from the ASA (May Be Required)

To access the IPS module CLI from the ASA, you can session from the ASA. For software modules, you can either session to the module (using Telnet) or create a virtual console session. A console session might be useful if the control plane is down and you cannot establish a Telnet session.

You may need to access the CLI if you are using multiple context mode and you need to set basic network settings using the CLI, or for troubleshooting.

Detailed Steps

Command	Purpose

Telnet session.
 	
For a physical module (for example, the ASA 5585-X):
session 1
 	
For a software module (for example, the ASA 5545-X):
session ips
 	
Example:
`hostname# session 1`
 	
`Opening command session with slot 1.`
`Connected to slot 1. Escape character sequence is 'CTRL-^X'.`
 	
`sensor login: cisco`
`Password: cisco`
	Accesses the module using Telnet. You are prompted for the username and password. The default username is **cisco**, and the default password is **cisco**.
	
	Note The first time you log in to the module, you are prompted to change the default password. Passwords must be at least eight characters long and cannot be a word in the dictionary.

Console session (software module only).
 	
session ips console
 	
Example:
`hostname# session ips console`
 	
`Establishing console session with slot 1`
`Opening console session with module ips.`
`Connected to module ips. Escape character sequence is 'CTRL-SHIFT-6 then x'.`
 	
`sensor login: cisco`
`Password: cisco`
	Accesses the module console. You are prompted for the username and password. The default username is **cisco**, and the default password is **cisco**.
	
	Note Do not use this command in conjunction with a terminal server where **Ctrl-Shift-6, x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6, x** is also the sequence to escape the IPS console and return to the ASA prompt. Therefore, if you try to exit the IPS console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the ASA, the IPS console session is still active; you can never exit to the ASA prompt. You must use a direct serial connection to return the console to the ASA prompt.
	
	Use the **session ips** command instead.

(ASA 5512-X through ASA 5555-X) Booting the Software Module

Your ASA typically ships with IPS module software present on Disk0. If the module is not running, or if you are adding the IPS module to an existing ASA, you must boot the module software. If you are unsure if the module is running, you will not see the IPS Basic COnfiguration screen when you run the Startup Wizard (see the "Configuring Basic IPS Module Network Settings" section on page 81-12).

Detailed Steps

Step 1 Do one of the following:

- New ASA with IPS pre-installed—To view the IPS module software filename in flash memory, choose **Tools > File Management**.

  For example, look for a filename like IPS-SSP_5512-K9-sys-1.1-a-7.1-4-E4.aip. Note the filename; you will need this filename later in the procedure.

- Existing ASA with new IPS installation—Download the IPS software from Cisco.com to your computer. If you have a Cisco.com login, you can obtain the software from the following website:

  http://www.cisco.com/cisco/software/navigator.html?mdfid=282164240

  Choose **Tools > File Management**, then choose **File Transfer > Between Local PC and Flash** to upload the new image to disk0. Note the filename; you will need this filename later in the procedure.

Step 2 Choose **Tools > Command Line Interface**.

Step 3 To set the IPS module software location in disk0, enter the following command and then click **Send**:

```
sw-module module ips recover configure image disk0:file_path
```

For example, using the filename in the example in Step 1, enter:

```
sw-module module ips recover configure image disk0:IPS-SSP_5512-K9-sys-1.1-a-7.1-4-E4.aip
```

Step 4 To install and load the IPS module software, enter the following command and then click **Send**:

```
sw-module module ips recover boot
```

Step 5 To check the progress of the image transfer and module restart process, enter the following command and then click **Send**:

```
show module ips details
```

The Status field in the output indicates the operational status of the module. A module operating normally shows a status of "Up." While the ASA transfers an application image to the module, the Status field in the output reads "Recover." When the ASA completes the image transfer and restarts the module, the newly transferred image is running.

Configuring Basic IPS Module Network Settings

- (ASA 5510 and Higher) Configuring Basic Network Settings, page 81-13
- (ASA 5505) Configuring Basic Network Settings, page 81-14

(ASA 5510 and Higher) Configuring Basic Network Settings

In single context mode, you can use the Startup Wizard in ASDM to configure basic IPS network configuration. These settings are saved to the IPS configuration, not the ASA configuration.

In multiple context mode, session to the module from the ASA and configure basic settings using the **setup** command.

Note (ASA 5512-X through ASA 5555-X) If you do not see the IPS Basic Configuration screen in your wizard, then the IPS module is not running. See the "(ASA 5512-X through ASA 5555-X) Booting the Software Module" section on page 81-12, and then repeat this procedure after you install the module.

Detailed Steps—Single Mode

Step 1 Choose **Wizards > Startup Wizard**.

Step 2 Click **Next** to advance through the initial screens until you reach the IPS Basic Configuration screen.

Step 3 In the Network Settings area, configure the following:

- IP Address—The management IP address. By default, the address is 192.168.1.2.
- Subnet Mask—The subnet mask for the management IP address.
- Gateway—The IP address of the upstream router. The IP address of the next hop router. See the "Connecting the ASA IPS Management Interface" section on page 81-8 to understand the requirements for your network. The default setting of the ASA management IP address will not work.
- HTTP Proxy Server—(Optional) The HTTP proxy server address. You can use a proxy server to download global correlation updates and other information instead of downloading over the Internet.
- HTTP Proxy Port—(Optional) The HTTP proxy server port.
- DNS Primary—(Optional) The primary DNS server address. If you are using a DNS server, you must configure at least one DNS server and it must be reachable for global correlation updates to be successful.

  For global correlation to function, you must have either a DNS server or an HTTP proxy server configured at all times. DNS resolution is supported only for accessing the global correlation update server.

Step 4 In the Management Access List area, enter an IP address and subnet mask for any hosts that are allowed to access the IPS management interface, and click **Add**. You can add multiple IP addresses.

Step 5 In the Cisco Account Password area, set the password for the username **cisco** and confirm it. The username **cisco** and this password are used for Telnet sessions from hosts specified by the management access list and when accessing the IPS module from ASDM (Configuration > IPS). By default, the password is **cisco**.

Step 6 In the Network Participation area, which you use to have the IPS module participate in SensorBase data sharing, click **Full**, **Partial**, or **Off**.

Detailed Steps—Multiple Mode Using the CLI

	Command	Purpose
Step 1	Session to the IPS module according to the "Sessioning to the Module from the ASA (May Be Required)" section on page 81-11.	

Step 2	`setup`
	
	Example:
	`sensor# setup`
		Runs the setup utility for initial configuration of the ASA IPS module. You are prompted for basic settings. For the default gateway, specify the IP address of the upstream router. See the "Connecting the ASA IPS Management Interface" section on page 81-8 to understand the requirements for your network. The default setting of the ASA management IP address will not work.

(ASA 5505) Configuring Basic Network Settings

An ASA IPS module on the ASA 5505 does not have any external interfaces. You can configure a VLAN to allow access to an internal IPS management IP address over the backplane. By default, VLAN 1 is enabled for IPS management. You can only assign one VLAN as the management VLAN. This section describes how to change the management VLAN and IP address if you do not want to use the default, and how to set other required network parameters.

Note Perform this configuration on the ASA 5505, not on the ASA IPS module.

Prerequisites

When you change the IPS VLAN and management address from the default, be sure to also configure the matching ASA VLAN and switch port(s) according to the procedures listed in Chapter 16, "Starting Interface Configuration (ASA 5505)." You must define and configure the VLAN for the ASA so the IPS management interface is accessible on the network.

Restrictions

Do not configure NAT for the management address if you intend to access it using ASDM. For initial setup with ASDM, you need to access the real address. After initial setup (where you set the password on the ASA IPS module), you can configure NAT and supply ASDM with the translated address for accessing the ASA IPS module.

Detailed Steps

Step 1 In ASDM, choose **Configuration > Device Setup > SSC Setup**.

Note The following settings are written to the ASA IPS module application configuration, not the ASA configuration.

Step 2 In the Management Interface area, set the following:

a. Choose the Interface VLAN from the drop-down list.

This setting allows you to manage the ASA IPS module using this VLAN.

b. Enter the IPS management IP address. Make sure this address is on the same subnet as the ASA VLAN IP address. For example, if you assigned 10.1.1.1 to the VLAN for the ASA, then assign another address on that network, such as 10.1.1.2, for the IPS management address. By default, the address is 192.168.1.2

c. Choose the subnet mask from the drop-down list.

d. Enter the default gateway IP address.

Set the gateway to be the ASA IP address for the management VLAN. By default, this IP address is 192.168.1.1.

Step 3 In the Management Access List area, enter the following:

a. Enter the IP address for the management host network.

b. Choose the subnet mask from the drop-down list.

c. Click **Add** to add these settings to the Allowed Hosts/Networks list.

Step 4 In the IPS Password area, do the following:

a. Enter the current password. The default password is **cisco**.

b. Enter the new password, and confirm the change.

Step 5 Click **Apply** to save the settings to the running configuration.

Step 6 To launch the IPS Startup Wizard, click the **Configure the IPS SSC module** link.

Configuring the Security Policy on the ASA IPS Module

This section describes how to configure the ASA IPS module application.

Detailed Steps

Step 1 Connect to ASDM using the *ASA* management IP address. See the "Starting ASDM" section on page 3-13.

Step 2 To access the IPS Device Manager (IDM) from ASDM, click **Configuration > IPS**.

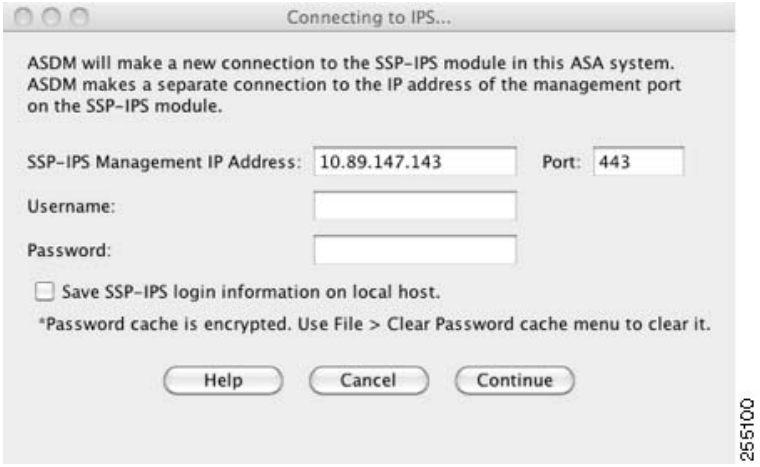

Step 3 Enter the IP address, username and password that you set in the "Configuring Basic IPS Module Network Settings" section on page 81-12, as well as the port. The default IP address and port is 192.168.1.2:443. The default username and password is **cisco** and **cisco**.

If the password to access IDM is lost, you can reset the password using ASDM. See the "Resetting the Password" section on page 81-22, for more information.

Step 4 To save the login information on your local PC, check the **Save IPS login information on local host** check box.

Step 5 Click **Continue**.

The Startup Wizard pane appears.

Step 6 Click **Launch Startup Wizard**. Complete the screens as prompted. For more information, see the IDM online help.

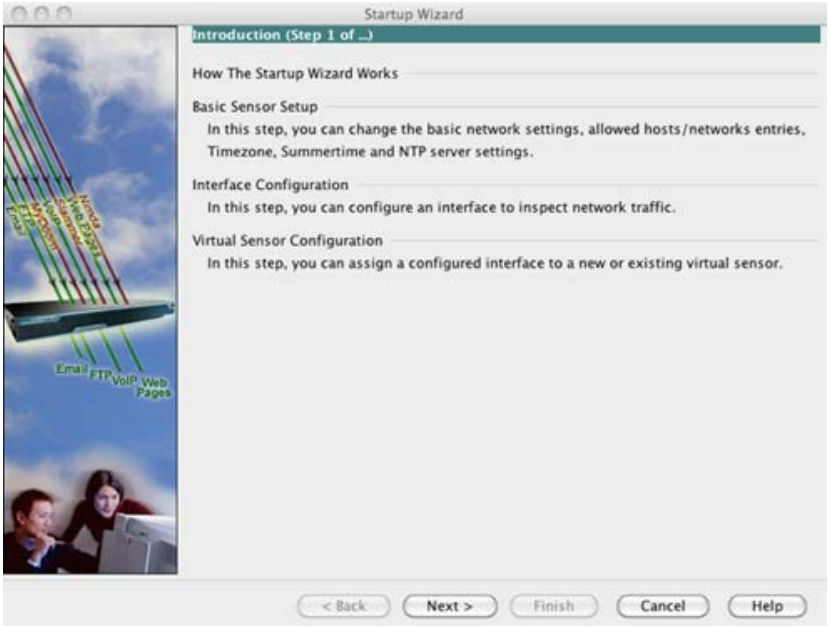

(ASA 5510 and higher) If you configure virtual sensors, you identify one of the sensors as the default. If the ASA series does not specify a virtual sensor name in its configuration, the default sensor is used.

What to Do Next

- For the ASA in multiple context mode, see the "Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher)" section on page 81-17.
- For the ASA in single context mode, see the "Diverting Traffic to the ASA IPS module" section on page 81-18.

Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher)

If the ASA is in multiple context mode, then you can assign one or more IPS virtual sensors to each context. Then, when you configure the context to send traffic to the ASA IPS module, you can specify a sensor that is assigned to the context; you cannot specify a sensor that you did not assign to the context. If you do not assign any sensors to a context, then the default sensor configured on the ASA IPS module is used. You can assign the same sensor to multiple contexts.

Note You do not need to be in multiple context mode to use virtual sensors; you can be in single mode and use different sensors for different traffic flows.

Prerequisites

For more information about configuring contexts, see the "Configuring Multiple Contexts" section on page 9-15.

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 Multiple Contexts" section on page 9-15.

Step 3 In the IPS Sensor Allocation area, click **Add**.

The IPS Sensor Selection dialog box appears.

Step 4 From the Sensor Name drop-down list, choose a sensor name from those configured on the ASA IPS module.

Step 5 (Optional) To assign a mapped name to the sensor, enter a value in the Mapped Sensor Name field.

This sensor name can be used within the context instead of the actual sensor name. If you do not specify a mapped name, the sensor name is used within the context. For security purposes, you might not want the context administrator to know which sensors are being used by the context. Or you might want to genericize the context configuration. For example, if you want all contexts to use sensors called "sensor1" and "sensor2," then you can map the "highsec" and "lowsec" sensors to sensor1 and sensor2 in context A, but map the "medsec" and "lowsec" sensors to sensor1 and sensor2 in context B.

Step 6 Click **OK** to return to the Edit Context dialog box.

Step 7 (Optional) To set one sensor as the default sensor for this context, from the Default Sensor drop-down list, choose a sensor name.

If you do not specify a sensor name when you configure IPS within the context configuration, the context uses this default sensor. You can only configure one default sensor per context. If you do not specify a sensor as the default, and the context configuration does not include a sensor name, then traffic uses the default sensor on the ASA IPS module.

Step 8 Repeat this procedure for each security context.

Step 9 Change to each context to configure the IPS security policy as described in "Diverting Traffic to the ASA IPS module" section on page 81-18.

What to Do Next

Change to each context to configure the IPS security policy as described in "Diverting Traffic to the ASA IPS module" section on page 81-18.

Diverting Traffic to the ASA IPS module

This section identifies traffic to divert from the ASA to the ASA IPS module.

Prerequisites

In multiple context mode, perform these steps in each context execution space. To change to a context, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Service Policy Rules**.

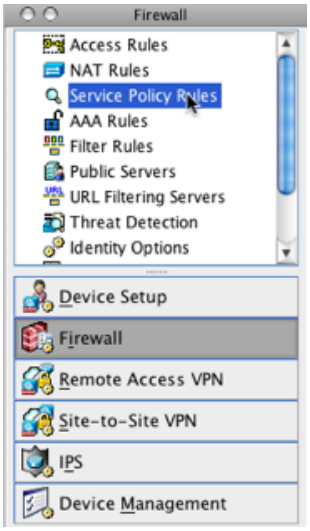

Step 2 Choose **Add > Add Service Policy Rule**. The Add Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3 Complete the Service Policy dialog box, and then the Traffic Classification Criteria dialog box as desired. See the ASDM online help for more information about these screens.

Step 4 Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 5 Click the **Intrusion Prevention** tab.

Step 6 Check the **Enable IPS for this traffic flow** check box.

Step 7 In the Mode area, click **Inline Mode** or **Promiscuous Mode**. See the "Operating Modes" section on page 81-3 for more information.

Step 8 In the If IPS Card Fails area, click **Permit traffic** or **Close traffic**. The Close traffic option sets the ASA to block all traffic if the ASA IPS module is unavailable. The Permit traffic option sets the ASA to allow all traffic through, uninspected, if the ASA IPS module is unavailable. For information about the IPS Sensor Selection area, see the ASDM online help.

Step 9 (ASA 5510 and higher) From the IPS Sensor to use drop-down list, choose a virtual sensor name.

If you use virtual sensors, you can specify a sensor name using this option. If you use multiple context mode on the ASA, you can only specify sensors that you assigned to the context (see the "Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher)" section on page 81-17). If you do not specify a sensor name, then the traffic uses the default sensor. In multiple context mode, you can specify a default sensor for the context. In single mode or if you do not specify a default sensor in multiple mode, the traffic uses the default sensor that is set on the ASA IPS module.

Step 10 Click **OK** and then **Apply**.

Step 11 Repeat this procedure to configure additional traffic flows as desired.

Managing the ASA IPS module

This section includes procedures that help you recover or troubleshoot the module and includes the following topics:

- Installing and Booting an Image on the Module, page 81-20
- Shutting Down the Module, page 81-22

- Uninstalling a Software Module Image, page 81-22
- Resetting the Password, page 81-22
- Reloading or Resetting the Module, page 81-23

Installing and Booting an Image on the Module

If the module suffers a failure, and the module application image cannot run, you can reinstall a new image on the module from a TFTP server (for a physical module), or from the local disk (software module).

Note Do not use the **upgrade** command within the module software to install the image.

Prerequisites

- Physical module—Be sure the TFTP server that you specify can transfer files up to 60 MB in size.

  Note This process can take approximately 15 minutes to complete, depending on your network and the size of the image.

- Software module—Copy the image to the ASA internal flash (disk0) before completing this procedure.

  Note Before you download the IPS software to disk0, make sure at least 50% of the flash memory is free. When you install IPS, IPS reserves 50% of the internal flash memory for its file system.

Detailed Steps

	Command	Purpose

Step 1	For a physical module (for example, the ASA 5585-X):
	hw-module module 1 recover configure
	For a software module (for example, the ASA 5545-X):
	sw-module module ips recover configure image disk0:*file_path*
	Example:
	`hostname# hw-module module 1 recover configure`
	`Image URL [tftp://127.0.0.1/myimage]: tftp://10.1.1.1/ids-newimg`
	`Port IP Address [127.0.0.2]: 10.1.2.10`
	`Port Mask [255.255.255.254]: 255.255.255.0`
	`Gateway IP Address [1.1.2.10]: 10.1.2.254`
	`VLAN ID [0]: 100`
		Specifies the location of the new image.
		For a physical module—This command prompts you for the URL for the TFTP server, the management interface IP address and netmask, gateway address, and VLAN ID (ASA 5505 only). These network parameters are configured in ROMMON; the network parameters you configured in the module application configuration are not available to ROMMON, so you must set them separately here.
		For a software module—Specify the location of the image on the local disk.
		You can view the recovery configuration using the **show module** {**1** | **ips**} **recover** command.
		In multiple context mode, enter this command in the system execution space.

Step 2	For a physical module: **hw-module module 1 recover boot** For a software module: **sw-module module ips recover boot** **Example:** `hostname# hw-module module 1 recover boot`	Installs and boots the IPS module software.
Step 3	For a physical module: **show module 1 details** For a software module: **show module ips details** **Example:** `hostname# show module 1 details`	Checks the progress of the image transfer and module restart process. The Status field in the output indicates the operational status of the module. A module operating normally shows a status of "Up." While the ASA transfers an application image to the module, the Status field in the output reads "Recover." When the ASA completes the image transfer and restarts the module, the newly transferred image is running.

Shutting Down the Module

Shutting down the module software prepares the module to be safely powered off without losing configuration data. **Note**: If you reload the ASA, the module is not automatically shut down, so we recommend shutting down the module before reloading the ASA. To gracefully shut down the module, perform the following steps at the ASA CLI.

Detailed Steps

Command	Purpose
For a physical module (for example, the ASA 5585-X): **hw-module module 1 shutdown** For a software module (for example, the ASA 5545-X): **sw-module module ips shutdown** **Example:** `hostname# hw-module module 1 shutdown`	Shuts down the module.

Uninstalling a Software Module Image

To uninstall a software module image and associated configuration, perform the following steps.

Detailed Steps

Command	Purpose

sw-module module ips uninstall
Example:
`hostname# sw-module module ips uninstall`
`Module ips will be uninstalled. This will completely remove the`
`disk image associated with the sw-module including any configuration`
`that existed within it.`
`Uninstall module <id>? [confirm]`
	Permanently uninstalls the software module image and associated configuration.

Resetting the Password

You can reset the module password to the default. The default password is **cisco**. After resetting the password, you should change it to a unique value using the module application.

Resetting the module password causes the module to reboot. Services are not available while the module is rebooting.

If you cannot connect to ASDM with the new password, restart ASDM and try to log in again. If you defined a new password and still have an existing password in ASDM that is different from the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

To reset the module password to the default of cisco, perform the following steps.

Detailed Steps

Step 1 From the ASDM menu bar, choose **Tools >** *module* **Password Reset**.

The Password Reset confirmation dialog box appears.

Step 2 Click **OK** to reset the password to the default.

A dialog box displays the success or failure of the password reset.

Step 3 Click **Close** to close the dialog box.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Detailed Steps

Command	Purpose
For a physical module (for example, the ASA 5585-X): **hw-module module 1 reload** For a software module (for example, the ASA 5545-X): **sw-module module ips reload** **Example:** `hostname# hw-module module 1 reload`	Reloads the module software.
For a physical module: **hw-module module 1 reset** For a software module: **sw-module module ips reset** **Example:** `hostname# hw-module module 1 reset`	Performs a reset, and then reloads the module.

Monitoring the ASA IPS module

See the "Intrusion Prevention Tab" section on page 4-28.

Feature History for the ASA IPS module

Table 81-2 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 81-2 Feature History for the ASA IPS module

Feature Name	Platform Releases	Feature Information
AIP SSM	7.0(1)	We introduced support for the AIP SSM for the ASA 5510, 5520, and 5540. The following screen was introduced: Configuration > Firewall > Service Policy Rules > Add/Edit Service Policy Rule > Intrusion Prevention.
Virtual sensors (ASA 5510 and higher)	8.0(2)	Virtual sensor support was introduced. Virtual sensors let you configure multiple security policies on the ASA IPS module. The following screen was modified: Context Management > Security Contexts > Edit Context.
AIP SSC for the ASA 5505	8.2(1)	We introduced support for the AIP SSC for the ASA 5505. The following screen was introduced: Configuration > Device Setup > SSC Setup.
Support for the ASA IPS SSP-10, -20, -40, and -60 for the ASA 5585-X	8.2(5)/ 8.4(2)	We introduced support for the ASA IPS SSP-10, -20, -40, and -60 for the ASA 5585-X. You can only install the ASA IPS SSP with a matching-level SSP; for example, SSP-10 and ASA IPS SSP-10. **Note** The ASA 5585-X is not supported in Version 8.3.

Table 81-2 Feature History for the ASA IPS module (continued)

Feature Name	Platform Releases	Feature Information
Support for Dual SSPs for SSP-40 and SSP-60	8.4(2)	For SSP-40 and SSP-60, you can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired. **Note** When using two SSPs in the chassis, VPN is not supported; note, however, that VPN has not been disabled. We did not modify any screens.
Support for the ASA IPS SSP for the ASA 5512-X through ASA 5555-X	8.6(1)	We introduced support for the ASA IPS SSP software module for the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X. We did not modify any screens.

Configuring the ASA CX Module

This chapter describes how to configure the ASA CX module that runs on the ASA. This chapter includes the following sections:

- Information About the ASA CX Module, page 82-1
- Licensing Requirements for the ASA CX Module, page 82-4
- Guidelines and Limitations, page 82-4
- Default Settings, page 82-5
- Configuring the ASA CX Module, page 82-5
- Managing the ASA CX Module, page 82-14
- Monitoring the ASA CX Module, page 82-15
- Troubleshooting the ASA CX Module, page 82-20
- Feature History for the ASA CX Module, page 82-21

Information About the ASA CX Module

The ASA CX module lets you enforce security based on the complete context of a situation. This context includes the identity of the user (who), the application or website that the user is trying to access (what), the origin of the access attempt (where), the time of the attempted access (when), and the properties of the device used for the access (how). With the ASA CX module, you can extract the full context of a flow and enforce granular policies such as permitting access to Facebook but denying access to games on Facebook or permitting finance employees access to a sensitive enterprise database but denying the same to other employees.

This section includes the following topics:

- How the ASA CX Module Works with the ASA, page 82-2
- Information About ASA CX Management, page 82-2
- Information About Authentication Proxy, page 82-3
- Information About VPN and the ASA CX Module, page 82-4
- Compatibility with ASA Features, page 82-4

How the ASA CX Module Works with the ASA

The ASA CX module runs a separate application from the ASA. The ASA CX module includes external management interface(s) so you can connect to the ASA CX module directly. Any data interfaces on the ASA CX module are used for ASA traffic only.

Traffic goes through the firewall checks before being forwarded to the ASA CX module. When you identify traffic for ASA CX inspection on the ASA, traffic flows through the ASA and the ASA CX module as follows:

1. Traffic enters the ASA.
2. Incoming VPN traffic is decrypted.
3. Firewall policies are applied.
4. Traffic is sent to the ASA CX module.
5. The ASA CX module applies its security policy to the traffic, and takes appropriate actions.
6. Valid traffic is sent back to the ASA; the ASA CX module might block some traffic according to its security policy, and that traffic is not passed on.
7. Outgoing VPN traffic is encrypted.
8. Traffic exits the ASA.

Figure 82-1 shows the traffic flow when using the ASA CX module. In this example, the ASA CX module automatically blocks traffic that is not allowed for a certain application. All other traffic is forwarded through the ASA.

Figure 82-1 ASA CX Module Traffic Flow in the ASA

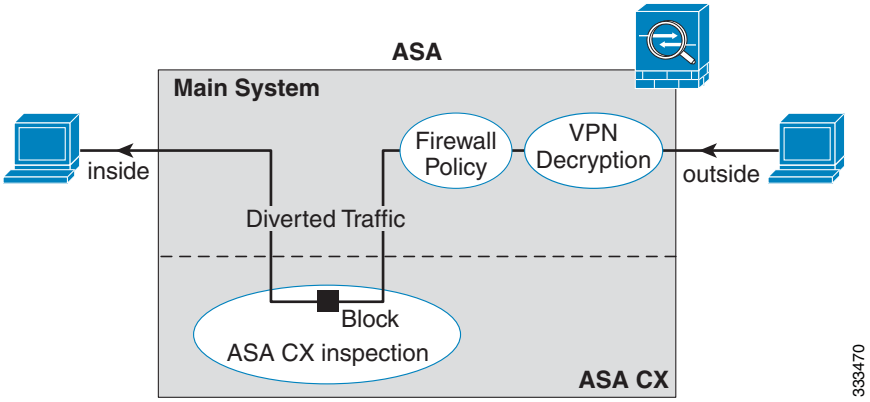

Note If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic orginiating on the non-ASA CX interface (the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because this feature is ingress-only.

Information About ASA CX Management

- Initial Configuration, page 82-3

- Policy Configuration and Management, page 82-3

Initial Configuration

For initial configuration, you must use the CLI on the ASA CX module to run the **setup** command and configure other optional settings.

To access the CLI, you can use the following methods:

- ASA CX console port—The ASA CX console port is a separate external console port.
- ASA CX Management 1/0 interface using SSH—You can connect to the default IP address (192.168.8.8), or you can use ASDM to change the management IP address and then connect using SSH. The ASA CX management interface is a separate external Gigabit Ethernet interface.

Note You cannot access the ASA CX module CLI over the ASA backplane using the **session** command.

Policy Configuration and Management

After you perform initial configuration, configure the ASA CX policy using Cisco Prime Security Manager (PRSM). Then configure the ASA policy for sending traffic to the ASA CX module using ASDM or the ASA CLI.

Note When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. Using PRSM lets you consolodate management to a single management system. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Information About Authentication Proxy

When the ASA CX needs to authenticate an HTTP user (to take advantage of identity policies), you must configure the ASA to act as an authentication proxy: the ASA CX module redirects authentication requests to the ASA interface IP address/proxy port. By default, the port is 885 (user configurable). Configure this feature as part of the service policy to divert traffic from the ASA to the ASA CX module. If you do not enable the authentication proxy, only passive authentication is available.

Note If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic orginiating on the non-ASA CX interface (the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because this feature is ingress-only.

Information About VPN and the ASA CX Module

The ASA includes VPN client and user authentication metadata when forwarding traffic to the ASA CX module, which allows the ASA CX module to include this information as part of its policy lookup criteria. The VPN metadata is sent only at VPN tunnel establishment time along with a type-length-value (TLV) containing the session ID. The ASA CX module caches the VPN metadata for each session. Each tunneled connection sends the session ID so the ASA CX module can look up that session's metadata.

Compatibility with ASA Features

The ASA includes many advanced application inspection features, including HTTP inspection. However, the ASA CX module provides more advanced HTTP inspection than the ASA provides, as well as additional features for other applications, including monitoring and controlling application usage.

To take full advantage of the ASA CX module features, see the following guidelines for traffic that you send to the ASA CX module:

- Do not configure ASA inspection on HTTP traffic.
- Do not configure Cloud Web Security (ScanSafe) inspection. If you configure both the ASA CX action and Cloud Web Security inspection for the same traffic, the ASA only performs the ASA CX action.
- Other application inspections on the ASA are compatible with the ASA CX module, including the default inspections.
- Do not enable the Mobile User Security (MUS) server; it is not compatible with the ASA CX module.
- Do not enable ASA clustering; they are not compatible; it is not compatible with the ASA CX module.
- If you enable failover, when the ASA fails over, any existing ASA CX flows are transferred to the new ASA, but the traffic is allowed through the ASA without being acted upon by the ASA CX module. Only new flows recieved by the new ASA are acted upon by the ASA CX module.

Licensing Requirements for the ASA CX Module

Model	License Requirement
All models	Base License.

The ASA CX module and PRSM require additional licenses. See the ASA CX documentation for more information.

Guidelines and Limitations

Context Mode Guidelines

Supported in single context mode only. Does not support multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support failover directly; when the ASA fails over, any existing ASA CX flows are transferred to the new ASA, but the traffic is allowed through the ASA without being inspected by the ASA CX.

ASA Clustering Guidelines

Does not support clustering.

IPv6 Guidelines

Supports IPv6.

Model Guidelines

Supported only on the ASA 5585-X. See the *Cisco ASA Compatibility Matrix* for more information:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

Additional Guidelines and Limitations

- See the "Compatibility with ASA Features" section on page 82-4.
- You cannot change the software type installed on the module; if you purchase an ASA CX module, you cannot later install other software on it.

Default Settings

Table 82-1 lists the default settings for the ASA CX module.

Table 82-1 Default Network Parameters

Parameters	Default
Management IP address	Management 1/0 192.168.8.8/24
Gateway	192.168.8.1/24
SSH Username	admin
Password	Admin123

Configuring the ASA CX Module

This section describes how to configure the ASA CX module and includes the following topics:

- Task Flow for the ASA CX Module, page 82-6
- Connecting the ASA CX Management Interface, page 82-7
- Configuring the ASA CX Management IP Address, page 82-8
- Configuring Basic ASA CX Settings at the ASA CX CLI, page 82-9

- Configuring the Security Policy on the ASA CX Module Using PRSM, page 82-11
- Redirecting Traffic to the ASA CX Module, page 82-12

Task Flow for the ASA CX Module

Configuring the ASA CX module is a process that includes configuration of the ASA CX security policy on the ASA CX module and then configuration of the ASA to send traffic to the ASA CX module. To configure the ASA CX module, perform the following steps:

Step 1 Cable the ASA CX management interfaces interface. See the "Connecting the ASA CX Management Interface" section on page 82-7.

Step 2 On the ASA, configure the ASA CX module management IP address for initial SSH access. See the "Configuring the ASA CX Management IP Address" section on page 82-8.

Step 3 On the ASA CX module, configure basic settings. You must use the CLI to configure these settings. See the "Configuring Basic ASA CX Settings at the ASA CX CLI" section on page 82-9.

Step 4 On the ASA CX module, configure the security policy using PRSM. See the "Configuring the Security Policy on the ASA CX Module Using PRSM" section on page 82-11.

Step 5 (Optional) On the ASA, configure the authentication proxy port. See the "(Optional) Configuring the Authentication Proxy Port" section on page 82-12.

Step 6 On the ASA, identify traffic to divert to the ASA CX module. See the "Redirecting Traffic to the ASA CX Module" section on page 82-12.

Note When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Connecting the ASA CX Management Interface

In addition to providing management access to the ASA CX module, the ASA CX management interface needs access to an HTTP proxy server or a DNS server and the Internet for signature updates and more. This section describes recommended network configurations. Your network may differ.

The ASA CX module includes a separate management interface from the ASA. For initial setup, you can connect with SSH to the ASA CX Management 1/0 interface using the default IP address (192.168.8.8/24). If you cannot use the default IP address, you can either use the console port or use ASDM to change the management IP address so you can use SSH.

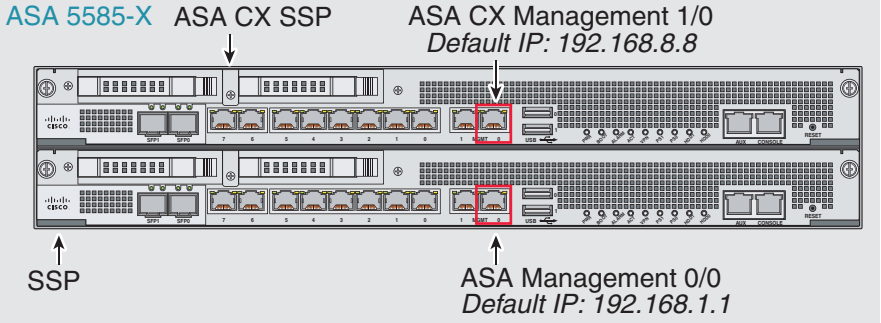

If you have an inside router

If you have an inside router, you can route between the management network, which can include both the ASA Management 0/0 and ASA CX Management 1/0 interfaces, and the ASA inside network for Internet access. Be sure to also add a route on the ASA to reach the Management network through the inside router.

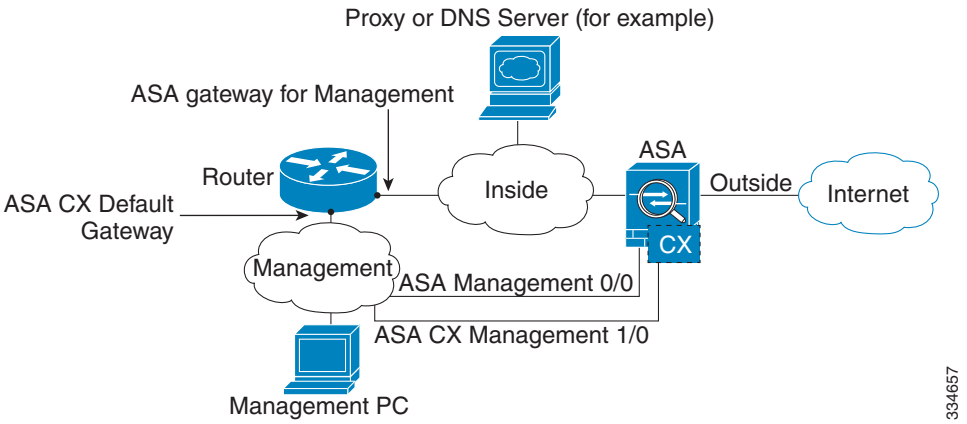

If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network, which would require an inside router to route between the networks. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. Because the ASA CX module is a separate device from the ASA, you can configure the ASA CX Management 1/0 address to be on the same network as the inside interface.

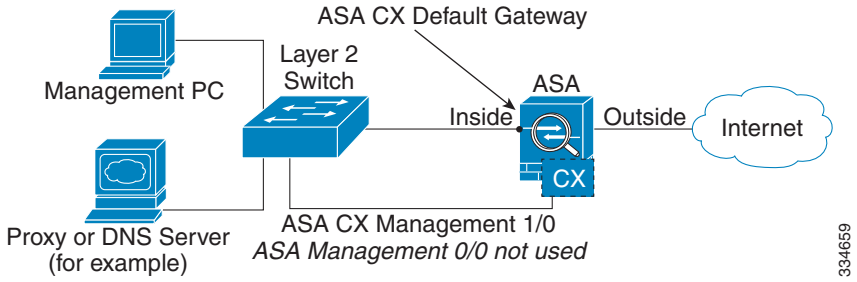

What to Do Next

- Configure the ASA CX management IP address. See the "Configuring the ASA CX Management IP Address" section on page 82-8.

Configuring the ASA CX Management IP Address

If you cannot use the default management IP address (192.168.8.8), then you can set the management IP address from the ASA. After you set the management IP address, you can access the ASA CX module using SSH to perform initial setup.

Detailed Steps

Step 1 In ASDM, choose **Wizards > Startup Wizard**.

Step 2 Click **Next** to advance through the initial screens until you reach the ASA CX Basic Configuration screen.

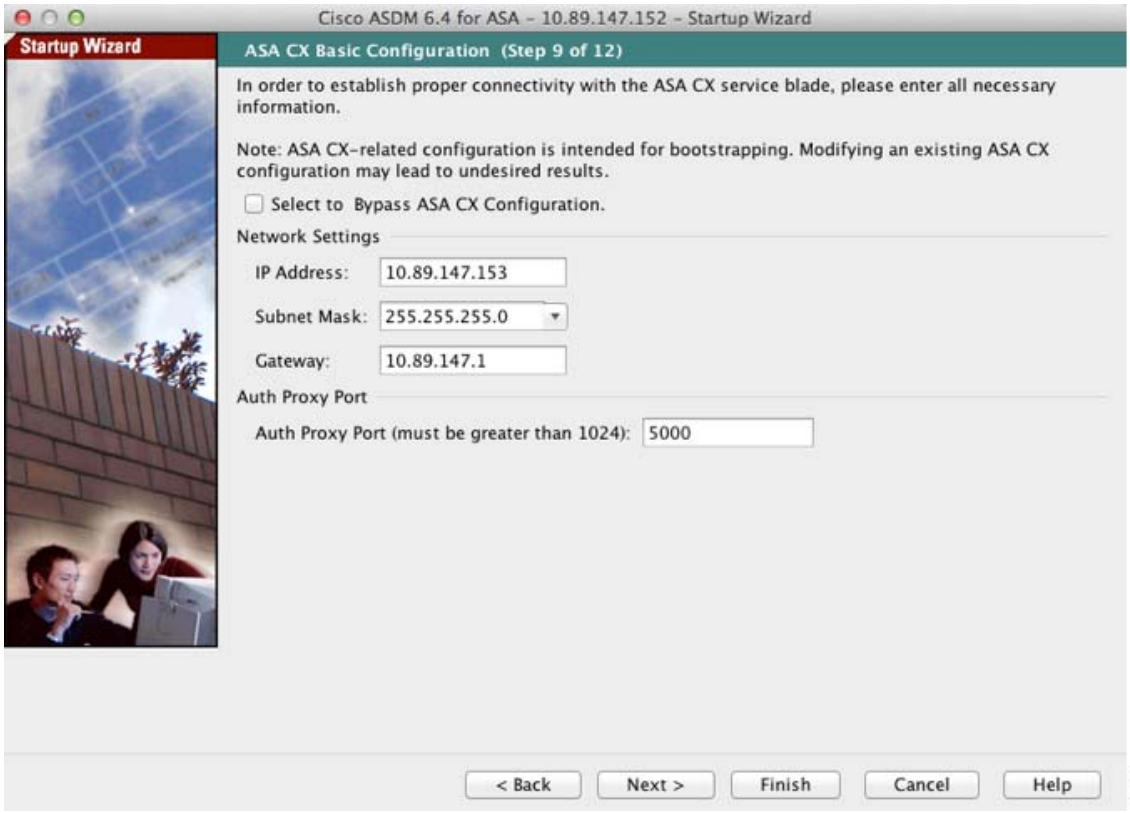

Step 3 Enter the new management IP address, subnet mask, and default gateway.

Step 4 (Optional) Change the Auth Proxy Port. You can set this later if desired. See the "(Optional) Configuring the Authentication Proxy Port" section on page 82-12 for more information.

Step 5 Click **Finish** to skip the remaining screens, or click **Next** to advance through the remaining screens and complete the wizard.

Configuring Basic ASA CX Settings at the ASA CX CLI

You must configure basic network settings and other parameters on the ASA CX module before you can configure your security policy.

Detailed Steps

Step 1 Connect to the ASA CX CLI:

- Using SSH to the ASA CX Management 1/0 interface—Log in with the username **admin** and the password **Admin123**. You will change the password as part of this procedure.
- Using the ASA CX console port.

Step 2 Enter the following command:

```
asacx> setup
```

Example:

```
asacx> setup
    Welcome to Cisco Prime Security Manager Setup
    [hit Ctrl-C to abort]
    Default values are inside [ ]
```

You are prompted through the setup wizard. The following example shows a typical path through the wizard; if you enter **Y** instead of **N** at a prompt, you will be able to configure some additional settings. This example shows how to configure both IPv4 and IPv6 static addresses. You can configure IPv6 stateless auto configuration by answering **N** when asked if you want to configure a static IPv6 address.

```
Enter a hostname [asacx]: asa-cx-host
Do you want to configure IPv4 address on management interface?(y/n) [Y]: Y
Do you want to enable DHCP for IPv4 address assignment on management interface?(y/n)[N]: N
Enter an IPv4 address [192.168.8.8]: 10.89.31.65
Enter the netmask [255.255.255.0]: 255.255.255.0
Enter the gateway [192.168.8.1]: 10.89.31.1
Do you want to configure static IPv6 address on management interface?(y/n) [N]: Y
Enter an IPv6 address: 2001:DB8:0:CD30::1234/64
Enter the gateway: 2001:DB8:0:CD30::1
Enter the primary DNS server IP address [ ]: 10.89.47.11
Do you want to configure Secondary DNS Server? (y/n) [N]: N
Do you want to configure Local Domain Name? (y/n) [N] Y
Enter the local domain name: example.com
Do you want to configure Search domains? (y/n) [N] Y
Enter the comma separated list for search domains: example.com
Do you want to enable the NTP service?(y/n) [N]: Y
Enter the NTP servers separated by commas: 1.ntp.example.com, 2.ntp.example.com
```

Step 3 After you complete the final prompt, you are presented with a summary of the settings. Look over the summary to verify that the values are correct, and enter **Y** to apply your changed configuration. Enter **N** to cancel your changes.

Example:

```
Apply the changes?(y,n) [Y]: Y
Configuration saved successfully!
Applying...
Done.
Generating self-signed certificate, the web server will be restarted after that
...
Done.
Press ENTER to continue...
asacx>
```

Note If you change the host name, the prompt does not show the new name until you log out and log back in.

Step 4 If you do not use NTP, configure the time settings. The default time zone is the UTC time zone. Use the **show time** command to see the current settings. You can use the following commands to change time settings:

```
asacx> config timezone
asacx> config time
```

Step 5 Change the admin password by entering the following command:

```
asacx> config passwd
```

Example:

```
asacx> config passwd
The password must be at least 8 characters long and must contain
at least one uppercase letter (A-Z), at least one lowercase letter
(a-z) and at least one digit (0-9).
Enter password: Farscape1
Confirm password: Farscape1
SUCCESS: Password changed for user admin
```

Step 6 Enter the **exit** command to log out.

Configuring the Security Policy on the ASA CX Module Using PRSM

This section describes how to launch PRSM to configure the ASA CX module application. For details on using PRSM to configure your ASA CX security policy, see the ASA CX user guide.

Detailed Steps

You can launch PRSM from your web browser, or you can launch it from ASDM.

- Launch PRSM from a web browser by enter the following URL:

  `https://`*`ASA_CX_management_IP`*

  Where the ASA CX management IP address is the one you set in the "Configuring Basic ASA CX Settings at the ASA CX CLI" section on page 82-9.

- Launch PRSM from ASDM by choosing **Home > ASA CX Status**, and clicking the **Connect to the ASA CX application** link.

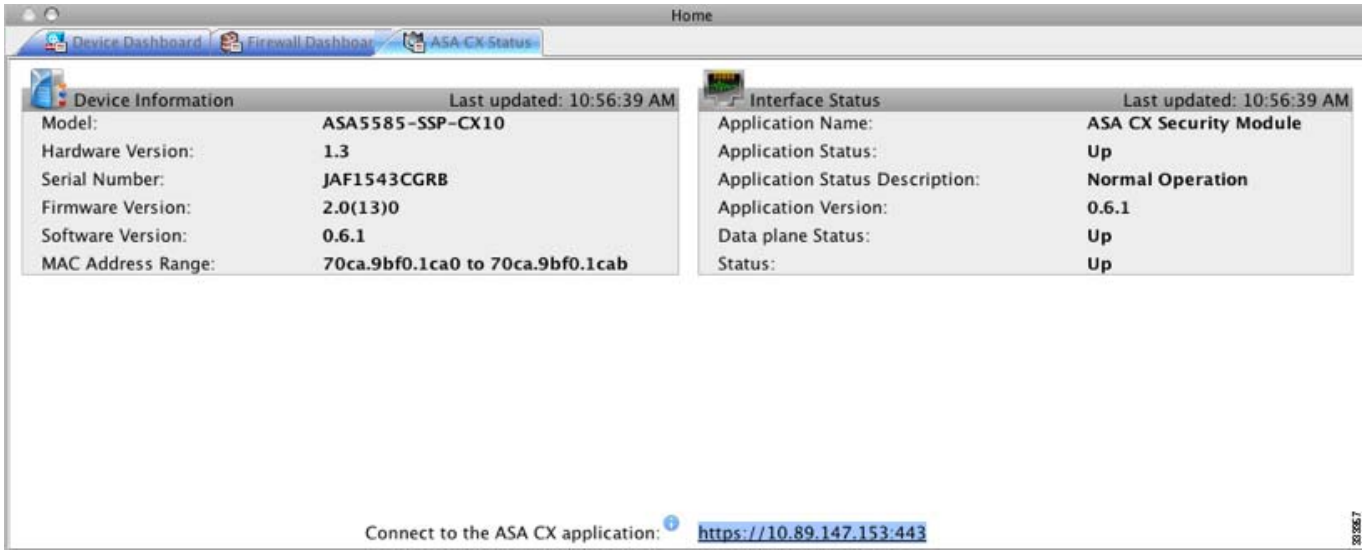

What to Do Next

- (Optional) Configure the authentication proxy port. See the "(Optional) Configuring the Authentication Proxy Port" section on page 82-12.

- Divert traffic to the ASA CX module. See the "Redirecting Traffic to the ASA CX Module" section on page 82-12.

(Optional) Configuring the Authentication Proxy Port

The default authentication port is 885. To change the authentication proxy port, perform the following steps. For more information about the authentication proxy, see the "Information About Authentication Proxy" section on page 82-3.

Note You can also set the port as part of the ASDM startup wizard. See the "Configuring Basic ASA CX Settings at the ASA CX CLI" section on page 82-9.

Detailed Steps

Step 1 In ASDM, choose **Configuration > Firewall > Advanced > ASA CX Auth Proxy**.

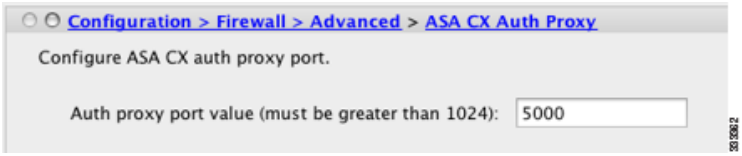

Step 2 Enter a port greater than 1024. The default is 885.

Step 3 Click **Apply**.

Redirecting Traffic to the ASA CX Module

This section identifies traffic to redirect from the ASA to the ASA CX module. Configure this policy on the ASA.

Note When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Prerequisites

If you enable the authentication proxy on the ASA using this procedure, be sure to also configure a directory realm for authentication on the ASA CX module. See the ASA CX user guide for more information.

Detailed Steps

Step 1 Choose **Configuration > Firewall > Service Policy Rules**.

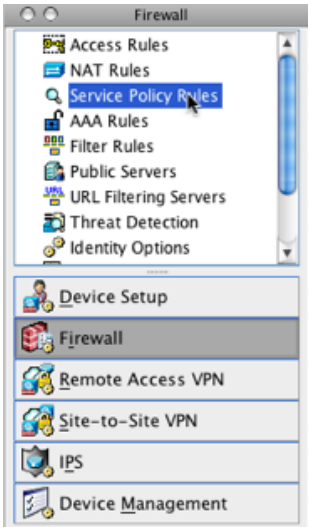

Step 2 Choose **Add > Add Service Policy Rule**. The Add Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3 Complete the Service Policy dialog box, and then the Traffic Classification Criteria dialog box as desired. See the ASDM online help for more information about these screens.

Step 4 Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 5 Click the **ASA CX Inspection** tab.

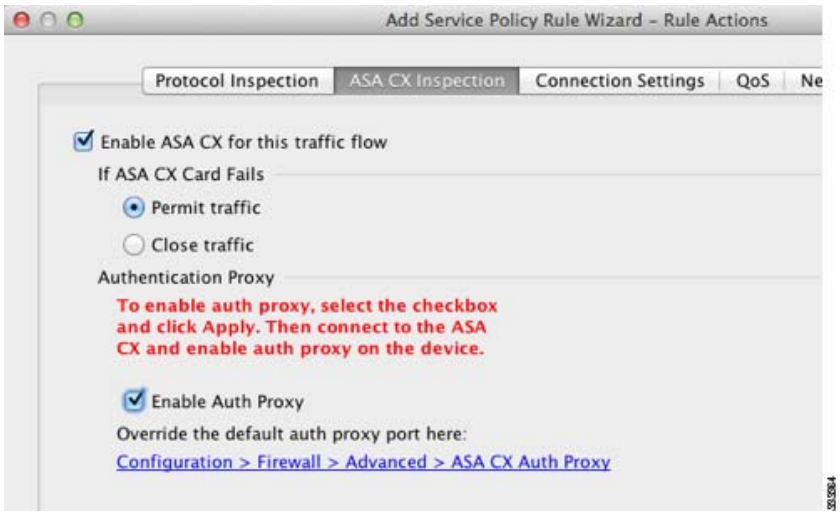

Step 6 Check the **Enable ASA CX for this traffic flow** check box.

Step 7 In the If ASA CX Card Fails area, click **Permit traffic** or **Close traffic**. The Close traffic option sets the ASA to block all traffic if the ASA CX module is unavailable. The Permit traffic option sets the ASA to allow all traffic through, uninspected, if the ASA CX module is unavailable.

Step 8 To enable the authentication proxy, which is required for active authentication, check the **Enable Auth Proxy** check box.

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

Step 10 Repeat this procedure to configure additional traffic flows as desired.

Managing the ASA CX Module

This section includes procedures that help you manage the module and includes the following topics:

- Resetting the Password, page 82-14
- Reloading or Resetting the Module, page 82-15
- Shutting Down the Module, page 82-15

Note You can install or upgrade your image from within the ASA CX module. See the ASA CX module documentaiton for more information.

Resetting the Password

You can reset the module password to the default. For the user **admin**, the default password is **Admin123**. After resetting the password, you should change it to a unique value using the module application.

Resetting the module password causes the module to reboot. Services are not available while the module is rebooting.

If you cannot connect to ASDM with the new password, restart ASDM and try to log in again. If you defined a new password and still have an existing password in ASDM that is different from the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

To reset the module password to the default of Admin123, perform the following steps.

Detailed Steps

Step 1 From the ASDM menu bar, choose **Tools > ASA CX Password Reset**.

The Password Reset confirmation dialog box appears.

Step 2 Click **OK** to reset the password to the default **Admin123**.

A dialog box displays the success or failure of the password reset.

Step 3 Click **Close** to close the dialog box.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Detailed Steps

Command	Purpose
hw-module module 1 reload **Example:** `hostname# hw-module module 1 reload`	Reloads the module software.
hw-module module 1 reset **Example:** `hostname# hw-module module 1 reset`	Performs a reset, and then reloads the module.

Shutting Down the Module

If you restart the ASA, the module is not automatically restarted. To shut down the module, perform the following steps at the ASA CLI.

Detailed Steps

Command	Purpose
hw-module module 1 shutdown **Example:** `hostname# hw-module module 1 shutdown`	Shuts down the module.

Monitoring the ASA CX Module

Use Tools > Command Line Interface to use monitoring commands.

- Showing Module Status, page 82-16
- Showing Module Statistics, page 82-16
- Monitoring Module Connections, page 82-17
- Capturing Module Traffic, page 82-20

- Problems with the Authentication Proxy, page 82-20

Note For ASA CX-related syslog messages, see the syslog messages guide. ASA CX syslog messages start with message number 429001.

Showing Module Status

See the "ASA CX Status Tab" section on page 4-30.

Showing Module Statistics

To show module statistics, enter the following command:

Command	Purpose
`show service-policy cxsc`	Displays the ASA CX statistics and status per service policy.

Examples

The following is sample output from the **show service-policy** command showing the ASA CX policy and the current statistics as well as the module status when the authentication proxy is disabled:

```
hostname# show service-policy cxsc
Global policy:
  Service-policy: global_policy
    Class-map: bypass
      CXSC: card status Up, mode fail-open, auth-proxy disabled
        packet input 2626422041, packet output 2626877967, drop 0, reset-drop 0, proxied 0
```

The following is sample output from the **show service-policy** command showing the ASA CX policy and the current statistics as well as the module status when the authentication proxy is enabled; in this case, the proxied counters also increment:

```
hostname# show service-policy cxsc
Global policy:
  Service-policy: pmap
    Class-map: class-default
      Default Queueing      Set connection policy: random-sequence-number disable
        drop 0
      CXSC: card status Up, mode fail-open, auth-proxy enabled
        packet input 7724, packet output 7701, drop 0, reset-drop 0, proxied 10
```

Monitoring Module Connections

To show connections through the ASA CX module, enter the one of the following commands:

Command	Purpose
show asp table classify domain cxsc	Shows the NP rules created to send traffic to the ASA CX module.
show asp table classify domain cxsc-auth-proxy	Shows the NP rules created for the authentication proxy for the ASA CX module.
show asp drop	Shows dropped packets. The following drop types are used: Frame Drops: • cxsc-bad-tlv-received—This occurs when ASA receives a packet from CXSC without a Policy ID TLV. This TLV must be present in non-control packets if it does not have the Standy Active bit set in the actions field. • cxsc-request—The frame was requested to be dropped by CXSC due a policy on CXSC whereby CXSC would set the actions to Deny Source, Deny Destination, or Deny Pkt. • cxsc-fail-close—The packet is dropped because the card is not up and the policy configured was 'fail-close' (rather than 'fail-open' which allows packets through even if the card was down). • cxsc-fail—The CXSC configuration was removed for an existing flow and we are not able to process it through CXSC it will be dropped. This should be very unlikely. • cxsc-malformed-packet—The packet from CXSC contains an invalid header. For instance, the header length may not be correct. Flow Drops: • cxsc-request—The CXSC requested to terminate the flow. The actions bit 0 is set. • reset-by-cxsc—The CXSC requested to terminate and reset the flow. The actions bit 1 is set. • cxsc-fail-close—The flow was terminated because the card is down and the configured policy was 'fail-close'.
show asp event dp-cp cxsc-msg	This output shows how many ASA CX module messages are on the dp-cp queue. Currently, only VPN queries from the ASA CX module are sent to dp-cp.
show conn	This command already shows if a connection is being forwarded to an module by displaying the ‘X - inspected by service module’ flag. Connections being forwarded to the ASA CX module will also display the ‘X’ flag.

Examples

The following is sample output from the **show asp table classify domain cxsc** command:

```
hostname# show asp table classify domain cxsc
Input Table
```

```
in  id=0x7ffedb4acf40, priority=50, domain=cxsc, deny=false
    hits=15485658, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0,
protocol=0
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
    input_ifc=outside, output_ifc=any
in  id=0x7ffedb4ad4a0, priority=50, domain=cxsc, deny=false
    hits=992053, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
    input_ifc=inside, output_ifc=any
in  id=0x7ffedb4ada00, priority=50, domain=cxsc, deny=false
    hits=0, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
    input_ifc=m, output_ifc=any

Output Table:

L2 - Output Table:

L2 - Input Table:

Last clearing of hits counters: Never
```

The following is sample output from the **show asp table classify domain cxsc-auth-proxy** command. For the first rule in the output below, the destination "port=2000" is the auth-proxy port configured by the **cxsc auth-proxy port 2000** command, and the destination "ip/id=192.168.0.100" is the ASA interface IP address.

```
hostname# show asp table classify domain cxsc-auth-proxy
Input Table
in  id=0x7ffed86cc470, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.0.100, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=inside, output_ifc=identity
in  id=0x7ffed86cce20, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=2.2.2.2, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=new2, output_ifc=identity
in  id=0x7ffed86cd7d0, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=172.23.58.52, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=mgmt, output_ifc=identity
in  id=0x7ffed86caa80, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.5.172, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=outside, output_ifc=identity
in  id=0x7ffed86cb3c0, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=::/0, port=0
        dst ip/id=fe80::5675:d0ff:fe5b:1102/128, port=2000
    input_ifc=outside, output_ifc=identity
in  id=0x7ffed742be10, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=::/0, port=0
    dst ip/id=1:1:1:1::10/128, port=2000
    input_ifc=outside, output_ifc=identity
```

```
Output Table:

L2 - Output Table:

L2 - Input Table:

Last clearing of hits counters: Never
```

The following is sample output from the **show asp drop** command. This output is just an example and lists all the possible reasons for a dropped frame or flow from the ASA CX module:

```
hostname# show asp drop
Frame drop:
  CXSC Module received packet with bad TLV's (cxsc-bad-tlv-received)           2
  CXSC Module requested drop (cxsc-request)                                    1
  CXSC card is down (cxsc-fail-close)                                          1
  CXSC config removed for flow (cxsc-fail)                                     3
  CXSC Module received malformed packet (cxsc-malformed-packet)                1

Last clearing: 18:12:58 UTC May 11 2012 by enable_15

Flow drop:
  Flow terminated by CXSC (cxsc-request)                                       2
  Flow reset by CXSC (reset-by-cxsc)                                           1
  CXSC fail-close (cxsc-fail-close)                                            1

Last clearing: 18:12:58 UTC May 11 2012 by enable_15
```

The following is sample output from the **show asp event dp-cp cxsc-msg** command:

```
hostname# show asp event dp-cp cxsc-msg
DP-CP EVENT QUEUE                   QUEUE-LEN  HIGH-WATER
Punt Event Queue                            0           5
Identity-Traffic Event Queue                0           0
General Event Queue                         0           4
Syslog Event Queue                          4          90
Non-Blocking Event Queue                    0           2
Midpath High Event Queue                    0          53
Midpath Norm Event Queue                 8074        8288
SRTP Event Queue                            0           0
HA Event Queue                              0           0
Threat-Detection Event Queue                0           3
ARP Event Queue                             0        2048
IDFW Event Queue                            0           0
CXSC Event Queue                            0           1
EVENT-TYPE          ALLOC ALLOC-FAIL ENQUEUED ENQ-FAIL  RETIRED 15SEC-RATE
cxsc-msg                1          0        1        0        1          0
```

The following is sample output from the **show conn detail** command:

```
hostname# show conn detail
0 in use, 105 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
```

```
TCP outside 208.80.152.2:80 inside 192.168.1.20:59928, idle 0:00:10, bytes 79174, flags
XUIO
```

Capturing Module Traffic

To configure and view packet captures for the ASA CX module, enter one of the following commands:

Command	Purpose
capture *name* **interface asa_dataplane**	Captures packets between ASA CX module and the ASA on the backplane.
copy capture	Copies the capture file to a server.
show capture	Shows the capture at the ASA console.

Note Captured packets contain an additional AFBP header that your PCAP viewer might not understand; be sure to use the appropriate plugin to view these packets.

Troubleshooting the ASA CX Module

- Problems with the Authentication Proxy, page 82-20

Problems with the Authentication Proxy

If you are having a problem using the authentication proxy feature, follow these steps to troubleshoot your configuration and connections:

1. Check your configurations.
   - On the ASA, check the output of the **show asp table classify domain cxsc-auth-proxy** command and make sure there are rules installed and that they are correct.
   - In PRSM, ensure the directory realm is created with the correct credentials and test the connection to make sure you can reach the authentication server; also ensure that a policy object or objects are configured for authentication.
2. Check the output of the **show service-policy cxsc** command to see if any packets were proxied.
3. Perform a packet capture on the backplane, and check to see if traffic is being redirected on the correct configured port. See the "Capturing Module Traffic" section on page 82-20. You can check the configured port using the **show running-config cxsc** command or the **show asp table classify domain cxsc-auth-proxy** command.

Note If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic orginiating on the non-ASA CX interface (the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because this feature is ingress-only.

Example 82-1 Make sure port 2000 is used consistently:

1. Check the authentication proxy port:

```
hostname# show running-config cxsc
cxsc auth-proxy port 2000
```

2. Check the authentication proxy rules:

```
hostname# show asp table classify domain cxsc-auth-proxy

Input Table
in  id=0x7ffed86cc470, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.0.100, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=inside, output_ifc=identity
```

3. In the packet captures, the redirect request should be going to destination port 2000.

Feature History for the ASA CX Module

Table 82-2 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 82-2 Feature History for the ASA CX Module

Feature Name	Platform Releases	Feature Information
ASA 5585-X support for the ASA CX SSP	8.4(4.1)	ASA CX module lets you enforce security based on the complete context of a situation. This context includes the identity of the user (who), the application or website that the user is trying to access (what), the origin of the access attempt (where), the time of the attempted access (when), and the properties of the device used for the access (how). With the ASA CX module, you can extract the full context of a flow and enforce granular policies such as permitting access to Facebook but denying access to games on Facebook or permitting finance employees access to a sensitive enterprise database but denying the same to other employees. We introduced the following screens: Home > ASA CX Status Wizards > Startup Wizard > ASA CX Basic Configuration Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA CX Inspection *This feature is not available in Version 8.6(1).*

Configuring the ASA CSC Module

This chapter describes how to configure the Content Security and Control (CSC) application that is installed in a CSC SSM in the ASA.

This chapter includes the following sections:

- Information About the CSC SSM, page 83-1
- Licensing Requirements for the CSC SSM, page 83-5
- Prerequisites for the CSC SSM, page 83-5
- Guidelines and Limitations, page 83-6
- Default Settings, page 83-6
- Configuring the CSC SSM, page 83-7
- CSC SSM Setup Wizard, page 83-10
- Using the CSC SSM GUI, page 83-20
- Monitoring the CSC SSM, page 83-24
- Troubleshooting the CSC Module, page 83-27
- Additional References, page 83-31
- Feature History for the CSC SSM, page 83-31

Information About the CSC SSM

Some ASA models support 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/HTTPS, POP3, and SMTP packets that you configure the ASA 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

Figure 83-1 shows the flow of traffic through an ASA that has the following:

- A CSC SSM installed and configured.
- A service policy that determines what traffic is diverted to the CSC SSM for scanning.

In this example, the client could be a network user who is accessing a website, downloading files from an FTP server, or retrieving mail from a POP3 server. SMTP scans differ in that you should configure the ASA to scan traffic sent from the outside to SMTP servers protected by the ASA.

Figure 83-1 Flow of Scanned Traffic with the CSC SSM

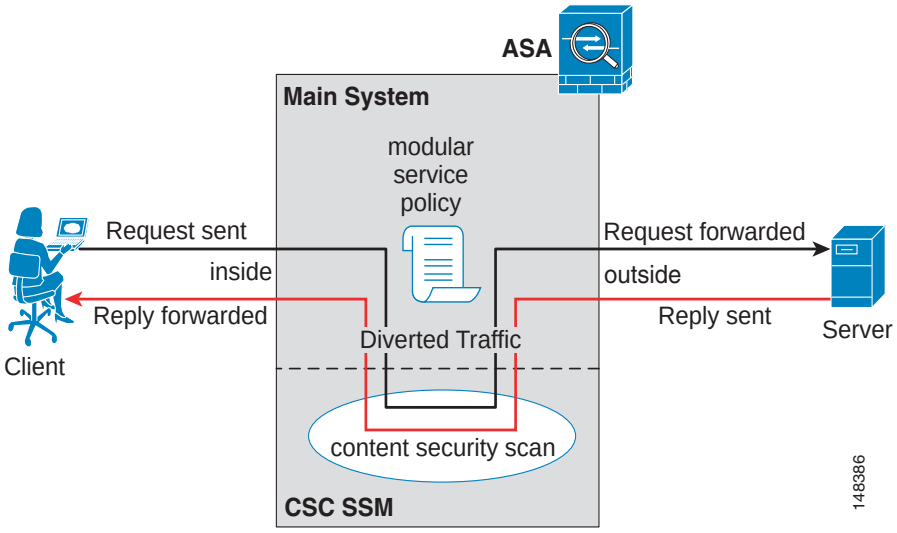

You use ASDM for system setup and monitoring of the CSC SSM. For advanced configuration of content security policies in the CSC SSM software, you access the web-based GUI for the CSC SSM by clicking links within ASDM. The CSC SSM GUI appears in a separate web browser window. To access the CSC SSM, you must enter the CSC SSM password. To use the CSC SSM GUI, see the *Cisco Content Security and Control SSM Administrator Guide*.

Note ASDM and the CSC SSM maintain separate passwords. You can configure their passwords to be identical; however, changing one of these two passwords does not affect the other password.

The connection between the host running ASDM and the ASA is made through a management port on the ASA. The connection to the CSC SSM GUI is made through the SSM management port. Because these two connections are required to manage the CSC SSM, any host running ASDM must be able to reach the IP address of both the ASA management port and the SSM management port.

Figure 83-2 shows an ASA with a CSC SSM that is connected to a dedicated management network. While use of a dedicated management network is not required, we recommend it. In this configuration, the following items are of particular interest:

- An HTTP proxy server is connected to the inside network and to the management network. This HTTP proxy server enables the CSC SSM to contact the Trend Micro Systems update server.
- The management port of the ASA is connected to the management network. To allow management of the ASA and the CSC SSM, hosts running ASDM must be connected to the management network.
- The management network includes an SMTP server for e-mail notifications for the CSC SSM and a syslog server to which the CSC SSM can send syslog messages.

Figure 83-2 CSC SSM Deployment with a Management Network

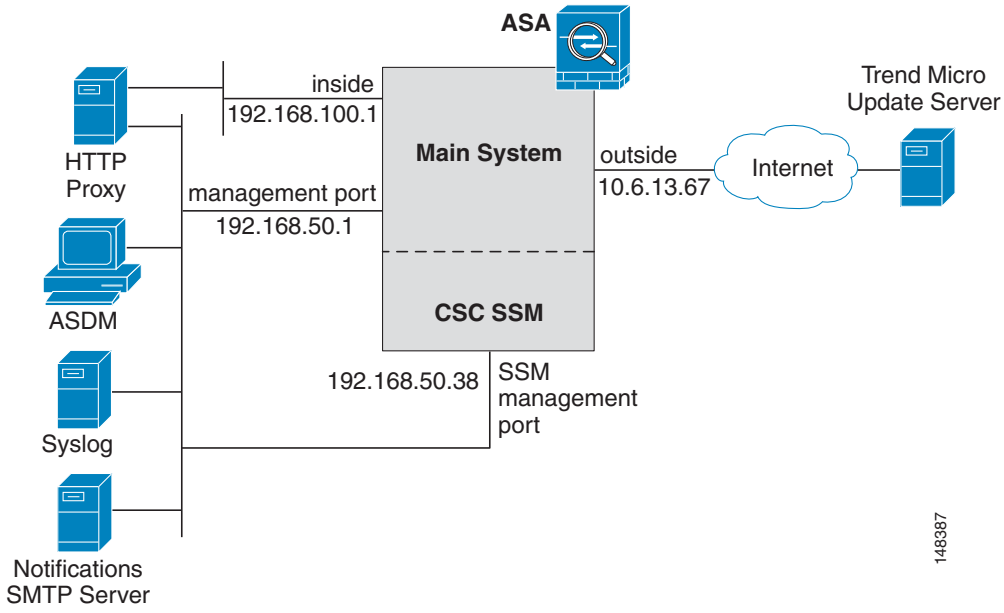

Determining What Traffic to Scan

The CSC SSM can scan FTP, HTTP/HTTPS, POP3, and SMTP traffic only when the destination port of the packet requesting the connection is the well-known port for the specified protocol. The CSC SSM can scan only the following connections:

- FTP connections opened to TCP port 21.
- HTTP connections opened to TCP port 80.
- HTTPS connections opened to TCP port 443.
- POP3 connections opened to TCP port 110.
- SMTP connections opened to TCP port 25.

You can choose to scan traffic for all of these protocols or any combination of them. For example, if you do not allow network users to receive POP3 e-mail, do not configure the ASA to divert POP3 traffic to the CSC SSM. Instead, block this traffic.

To maximize performance of the ASA and the CSC SSM, divert only the traffic to the CSC SSM that you want the CSC SSM to scan. Diverting traffic that you do not want scanned, such as traffic between a trusted source and destination, can adversely affect network performance.

Note When traffic is first classified for CSC inspection, it is flow-based. If traffic is part of a pre-existing connection, the traffic goes directly to the service policy set for that connection.

You can apply service policies that include CSC scanning globally or to specific interfaces; therefore, you can choose to enable CSC scans globally or for specific interfaces. For more information, see the "Determining Service Policy Rule Actions for CSC Scanning" section on page 83-9.

Based on the configuration shown in Figure 83-3, configure the ASA to divert to the CSC SSM only requests from clients on the inside network for HTTP, FTP, and POP3 connections to the outside network, and incoming SMTP connections from outside hosts to the mail server on the DMZ network. Exclude from scanning HTTP requests from the inside network to the web server on the DMZ network.

Figure 83-3 Common Network Configuration for CSC SSM Scanning

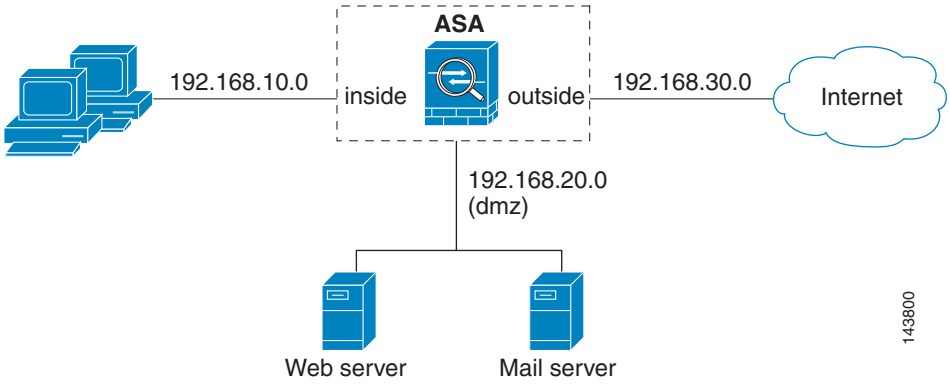

There are many ways you could configure the ASA to identify the traffic that you want to scan. One approach is to define two service policies: one on the inside interface and the other on the outside interface, each with access lists that match traffic to be scanned.

Figure 83-4 shows service policy rules that select only the traffic that the ASA should scan.

Figure 83-4 Optimized Traffic Selection for CSC Scans

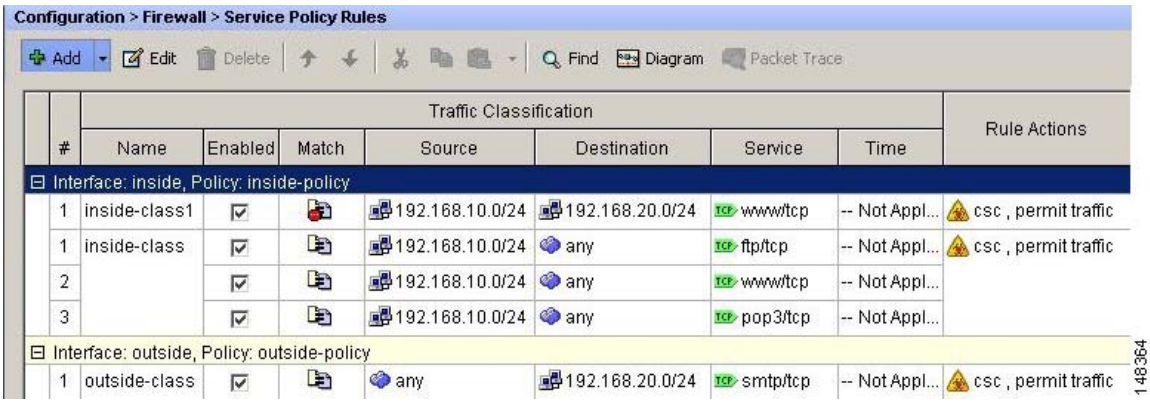

In the inside-policy, the first class, inside-class1, ensures that the ASA does not scan HTTP traffic between the inside network and the DMZ network. The Match column indicates this setting by displaying the “Do not match” icon. This setting does not mean the ASA blocks traffic sent from the 192.168.10.0 network to TCP port 80 on the 192.168.20.0 network. Instead, this setting exempts the traffic from being matched by the service policy applied to the inside interface, which prevents the ASA from sending the traffic to the CSC SSM.

The second class of the inside-policy, inside-class matches FTP, HTTP, and POP3 traffic between the inside network and any destination. HTTP connections to the DMZ network are exempted because of the inside-class1 setting. As previously mentioned, policies that apply CSC scanning to a specific interface affect both incoming and outgoing traffic, but by specifying 192.168.10.0 as the source network, inside-class1 matches only connections initiated by the hosts on the inside network.

In the outside-policy, outside-class matches SMTP traffic from any outside source to the DMZ network. This setting protects the SMTP server and inside users who download e-mail from the SMTP server on the DMZ network, without having to scan connections from SMTP clients to the server.

If the web server on the DMZ network receives files uploaded by HTTP from external hosts, you can add a rule to the outside policy that matches HTTP traffic from any source to the DMZ network. Because the policy is applied to the outside interface, the rule would only match connections from HTTP clients outside the ASA.

Licensing Requirements for the CSC SSM

Model	License Requirement
ASA 5510	• Base License—Supports SMTP virus scanning, POP3 virus scanning and content filtering, web mail virus scanning, HTTP file blocking, FTP virus scanning and file blocking, logging, and automatic updates. Supports two contexts. *Optional licenses: 5 contexts.* • Security Plus License—Supports the Base license features, plus SMTP anti-spam, SMTP content filtering, POP3 anti-spam, URL blocking, and URL filtering. Supports two contexts. *Optional license: 5 contexts.*
ASA 5520	Base License—Supports all features. Supports two contexts. *Optional licenses: 5, 10, or 20 contexts.*
ASA 5540	Base License—Supports all features. Supports two contexts. *Optional licenses: 5, 10, 20, or 50 contexts.*
All other models	No support.

Prerequisites for the CSC SSM

The CSC SSM has the following prerequisites:

- A CSC SSM card must be installed in the ASA.
- 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.
- E-mail address(es) for product license renewal 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 ASA 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.

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 ASA in which the CSC SSM is installed fails. When the standby ASA 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. Not supported on the ASA 5580 and the ASA 5585-X.

Additional Guidelines

You cannot change the software type installed on the module; if you purchase a CSC module, you cannot later install IPS software on it.

Default Settings

Table 83-1 lists the default settings for the CSC SSM.

Table 83-1 Default CSC SSM Parameters

Parameter	Default
FTP inspection on the ASA	Enabled
All features included in the license(s) that you have purchased	Enabled

Configuring the CSC SSM

This section describes how to configure the CSC SSM and includes the following topics:

- Before Configuring the CSC SSM, page 83-7
- Connecting to the CSC SSM, page 83-8
- Determining Service Policy Rule Actions for CSC Scanning, page 83-9

Before Configuring the CSC SSM

Before configuring the ASA and the CSC SSM, perform the following steps:

Step 1 If the CSC SSM did not come preinstalled in a Cisco ASA, install it and connect a network cable to the management port of the SSM. For assistance with installation and connecting the SSM, see the *Cisco ASA 5500 Series Quick Start Guide*.

The management port of the CSC SSM must be connected to your network to allow management of and automatic updates to the CSC SSM software. Additionally, the CSC SSM uses the management port for e-mail notifications and syslog messages.

Step 2 You should have received a Product Authorization Key (PAK) with the CSC SSM. Use the PAK to register the CSC SSM at the following URL.

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

After you register, you receive activation keys by e-mail. The activation keys are required before you can complete Step 6.

Step 3 Obtain the following information for use in Step 6:

- Activation keys
- 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 SMTP server IP address and port number for e-mail notifications
- E-mail address(es) for product license renewal notifications
- IP addresses of hosts or networks that are allowed to manage the CSC SSM
- Password for the CSC SSM

Step 4 In a web browser, access ASDM for the ASA in which the CSC SSM is installed.

Note If you are accessing ASDM for the first time, see the "Additional References" section on page 83-31.

For more information about enabling ASDM access, see the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 52-1.

Step 5 Verify time settings on the ASA. Time setting accuracy is important for logging of security events and for automatic updates of CSC SSM software. Do one of the following:

- If you manually control time settings, verify the clock settings, including time zone. Choose **Configuration > Properties > Device Administration > Clock**.
- If you are using NTP, verify the NTP configuration. Choose **Configuration > Properties > Device Administration > NTP**.

Step 6 Open ASDM.

Step 7 Connect to and log in to the CSC SSM. For instructions, see the "Connecting to the CSC SSM" section on page 83-8.

Step 8 Run the CSC Setup Wizard.

- To access the CSC Setup Wizard, choose **Configuration > Trend Micro Content Security > CSC Setup > Wizard Setup > Launch Setup Wizard**.
- If you are rerunning the CSC Setup Wizard, perform the same step listed in the previous bullet.

The CSC Setup Wizard appears.

Step 9 Complete the CSC Setup Wizard, which includes configuration of service policies to divert traffic that you want scanned to the CSC SSM.

Note If you create a global service policy to divert traffic for CSC scans, all traffic (inbound and outbound) for the supported protocols is scanned. To maximize performance of the ASA and the CSC SSM, scan traffic only from untrusted sources.

Step 10 To reduce the load on the CSC SSM, configure the service policy rules that send packets to the CSC SSM to support only HTTP/HTTPS, SMTP, POP3, or FTP traffic. For instructions, see the "Determining Service Policy Rule Actions for CSC Scanning" section on page 83-9.

Step 11 (Optional) Review the default content security policies in the CSC SSM GUI, which are suitable for most implementations. You review the content security policies by viewing the enabled features in the CSC SSM GUI. For the availability of features, see the "Licensing Requirements for the CSC SSM" section on page 83-5. For the default settings, see the "Default Settings" section on page 83-6.

What to Do Next

See the "Connecting to the CSC SSM" section on page 83-8.

Connecting to the CSC SSM

With each session you start in ASDM, the first time you access features related to the CSC SSM, you must specify the management IP address and provide the password for the CSC SSM. After you successfully connect to the CSC SSM, you are not prompted again for the management IP address and password. If you start a new ASDM session, the connection to the CSC SSM is reset and you must specify the IP address and the CSC SSM password again. The connection to the CSC SSM is also reset if you change the time zone on the ASA.

Note The CSC SSM has a password that is maintained separately from the ASDM password. You can configure the two passwords to be identical, but changing the CSC SSM password does not affect the ASDM password.

To connect to the CSC SSM, perform the following steps:

Step 1 In the ASDM main application window, click the **Content Security** tab.

Step 2 In the Connecting to CSC dialog box, click one of the following radio buttons:

- To connect to the IP address of the management port on the SSM, click **Management IP Address**. ASDM automatically detects the IP address for the SSM in the ASA. If this detection fails, you can specify the management IP address manually.
- To connect to an alternate IP address or hostname on the SSM, click **Other IP Address or Hostname**.

Step 3 Enter the port number in the Port field, and then click **Continue**.

Step 4 In the CSC Password field, type your CSC password, and then click **OK**.

Note If you have not completed the CSC Setup Wizard (choose **Configuration > Trend Micro Content Security > CSC Setup > Wizard Setup**), complete the configuration in the CSC Setup Wizard, which includes changing the default password, "cisco."

For ten minutes after you have entered the password, you do not need to reenter the CSC SSM password to access other parts of the CSC SSM GUI.

Step 5 To access the CSC SSM GUI, choose **Configuration > Trend Micro Content Security**, and then click one of the following tabs: **Web**, **Mail**, **File Transfer**, or **Updates**.

What to Do Next

See the "Determining Service Policy Rule Actions for CSC Scanning" section on page 83-9.

Determining Service Policy Rule Actions for CSC Scanning

The CSC SSM scans only HTTP/HTTPS, SMTP, POP3, and FTP traffic. If your service policy includes traffic that supports other protocols in addition to these four, packets for other protocols are passed through the CSC SSM without being scanned. You should configure the service policy rules that send packets to the CSC SSM to support only HTTP/HTTPS, SMTP, POP3, or FTP traffic.

The CSC Scan tab in the Add Service Policy Rule Wizard lets you determine whether or not the CSC SSM scans traffic identified by the current traffic class. This tab appears only if a CSC SSM is installed in the ASA.

To configure service policy rules for CSC scanning, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration > Firewall > Service Policy Rules**.

Step 2 On the toolbar, click **Add**.

The Add Service Policy Rule Wizard screen appears.

Step 3 Click the **Global - applies to all interfaces** option, and then click **Next**.

The Traffic Classification Criteria screen appears.

Step 4 Click the **Create a new traffic class** option, type a name for the traffic class in the adjacent field, check the **Any traffic** check box, and then click **Next**.

The Rule Actions screen appears.

Step 5 Click the **CSC Scan** tab, and then check the **Enable CSC scan for this traffic flow** check box.

Step 6 Choose whether the ASA should permit or deny selected traffic to pass if the CSC SSM is unavailable by making the applicable selection in the area labeled: If CSC card fails, then. When this check box is checked, the other parameters on this tab become active.

Step 7 In the If CSC card fails area, if the CSC SSM becomes inoperable, choose one of the following actions:

- To allow traffic, check the **Permit traffic** check box.
- To block traffic, check the **Close traffic** check box.

Step 8 Click **Finish**.

The new service policy rule appears in the Service Policy Rules pane.

Step 9 Click **Apply**.

The ASA begins diverting traffic to the CSC SSM, which performs the content security scans that have been enabled according to the license that you purchased.

CSC SSM Setup Wizard

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.

To start the CSC Setup Wizard, click **Launch Setup Wizard**.

This section includes the following topics:

- Activation/License, page 83-11
- IP Configuration, page 83-11
- Host/Notification Settings, page 83-12
- Management Access Host/Networks, page 83-13
- Password, page 83-13
- Restoring the Default Password, page 83-14
- Wizard Setup, page 83-15

Activation/License

The Activation/License pane lets you review or renew activation codes for the CSC SSM Basic 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 Choose **Configuration > Trend Micro Content Security > CSC Setup > Activation/License.**

Step 2 The Activation/License pane shows the following display-only information for the Basic 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 Basic 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. The Basic License includes anti-virus, anti-spyware, and file blocking. The Plus License includes anti-spam, anti-phishing, content filtering, URL blocking and filtering, and web reputation.

Step 3 To review license status or renew your license, click the link provided.

Step 4 To go to the CSC home pane in ASDM, click the link provided.

What to Do Next

See the "IP Configuration" section on page 83-11.

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 Choose **Configuration > Trend Micro Content Security > CSC Setup > IP Configuration**.

Step 2 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 3 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, if configured.

Step 4 (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, if configured.
- Enter the listening port of the proxy server, if configured.

What to Do Next

See the "Host/Notification Settings" section on page 83-12.

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 Choose **Configuration > Trend Micro Content Security > CSC Setup > Host/Notification Settings**.

Step 2 In the Host and Domain Names area, set the hostname and domain name of the CSC SSM.

Step 3 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 4 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.
- The e-mail address(es) for the product license renewal to which notification e-mails should be sent. Separate multiple e-mail addresses with semicolons. The maximum number of characters allowed for e-mail addresses is 1024. Make sure that the specified e-mail addresses are valid.

What to Do Next

See the "Management Access Host/Networks" section on page 83-13.

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 Choose **Configuration > Trend Micro Content Security > CSC Setup > Management Access Host/Networks**.

Step 2 Enter the IP address of a host or network that you want to add to the Selected Hosts/Network list.

Step 3 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 4 To add the host or network that you specified in the IP Address field in the Selected Hosts/Networks list, click **Add**.

The Selected Hosts/Networks table lists the IP addresses of networks and hosts whose connection to the CSC SSM you have added.

Step 5 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 83-13.

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 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

If required, see the "Restoring the Default Password" section on page 83-14.

See the "Wizard Setup" section on page 83-15.

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 ASA CLI by entering the **session** command. For more information, see the *Cisco Content Security and Control 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 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 ASA 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 83-13.

Wizard Setup

The Wizard Setup screen lets you start the CSC Setup Wizard. To start the CSC Setup Wizard, click **Launch Setup Wizard**. To access the Wizard Setup screen, choose **Configuration > Trend Micro Content Security > CSC Setup > Wizard Setup**.

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 83-15
- CSC Setup Wizard IP Configuration, page 83-16
- CSC Setup Wizard Host Configuration, page 83-16
- CSC Setup Wizard Management Access Configuration, page 83-17
- CSC Setup Wizard Password Configuration, page 83-17
- CSC Setup Wizard Traffic Selection for CSC Scan, page 83-17
- CSC Setup Wizard Summary, page 83-19

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 Basic License appears. The Basic 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 83-16.

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 Host Configuration" section on page 83-16.

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 SMTP server.
- The port to which the SMTP server listens.
- The e-mail address(es) for the product license renewal notification.

What to Do Next

See the "CSC Setup Wizard Management Access Configuration" section on page 83-17.

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 83-17.

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 83-17.

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 83-18.
- To modify additional traffic details for CSC scanning, click **Edit**. For more information, see "Specifying Traffic for CSC Scanning" section on page 83-18.
- 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 83-19.

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.
  - E-mail address(es) for product licensing renewal notifications.
- The settings that you made in 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 83-20.

Using the CSC SSM GUI

This section describes how to configure features using the CSC SSM GUI, and includes the following topics:

- Web, page 83-20
- Mail, page 83-21
- SMTP Tab, page 83-21
- POP3 Tab, page 83-22
- File Transfer, page 83-22
- Updates, page 83-23

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 83-21.

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 83-21
- POP3 Tab, page 83-22

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.

Step 7 The Global Approved List area is display-only and shows whether or not the SMTP global approved list feature is enabled on the CSC SSM. Click **Configure Global Approved List** to open a screen for configuring SMTP global approved list settings 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.

Step 5 The Global Approved List area is display-only and shows whether or not the POP3 global approved list feature is enabled on the CSC SSM. Click **Configure Global Approved List** to open a screen for configuring POP3 global approved list settings on the CSC SSM.

What to Do Next

See the "File Transfer" section on page 83-22.

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 83-23.

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 ASA, the SSM is not automatically restarted. For more information, see the "Managing SSMs and SSCs" section in the CLI configuration guide.

What to Do Next

See the "Monitoring the CSC SSM" section on page 83-24.

Monitoring the CSC SSM

ASDM lets you monitor the CSC SSM statistics as well as CSC SSM-related features.

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.

This section includes the following topics:

- Threats, page 83-24
- Live Security Events, page 83-25
- Live Security Events Log, page 83-25
- Software Updates, page 83-26
- Resource Graphs, page 83-27

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.

What to Do Next

See the "Live Security Events" section on page 83-25.

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.

What to Do Next

See the "Live Security Events Log" section on page 83-25.

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 scrolling of the Latest CSC Security Events pane, click **Pause**. To resume scrolling of the Latest CSC Security Events pane, click **Resume**.

Step 5 To save the log to a file on your PC, click **Save**.

Step 6 To clear the list of messages shown, click **Clear Display**.

Step 7 To close the pane and return to the previous one, click **Close**.

What to Do Next

See the "Software Updates" section on page 83-26.

Software Updates

To view information about CSC SSM software updates, 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.

What to Do Next

See the "CSC CPU" section on page 83-27.

Resource Graphs

The ASA lets you monitor CSC SSM status, including CPU resources and memory usage. This section includes the following topics:

- CSC CPU, page 83-27
- CSC Memory, page 83-27

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 83-24.

What to Do Next

See the "CSC Memory" section on page 83-27.

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 83-24.

Troubleshooting the CSC Module

This section includes procedures that help you recover or troubleshoot the module and includes the following topics:

- Installing an Image on the Module, page 83-28

- Resetting the Password, page 83-29
- Reloading or Resetting the Module, page 83-30
- Shutting Down the Module, page 83-30

Note This section covers all ASA module types; follow the steps appropriate for your module.

Installing an Image on the Module

If the module suffers a failure, and the module application image cannot run, you can reinstall a new image on the module from a TFTP server.

Note Do not use the **upgrade** command within the module software to install the image.

Prerequisites

Be sure the TFTP server that you specify can transfer files up to 60 MB in size.

Note This process can take approximately 15 minutes to complete, depending on your network and the size of the image.

Detailed Steps

	Command	Purpose

Step 1	**hw-module module 1 recover configure**
	
	Example:
	`hostname# hw-module module 1 recover configure`
	`Image URL [tftp://127.0.0.1/myimage]: tftp://10.1.1.1/ids-newimg`
	`Port IP Address [127.0.0.2]: 10.1.2.10`
	`Port Mask [255.255.255.254]: 255.255.255.0`
	`Gateway IP Address [1.1.2.10]: 10.1.2.254`
	`VLAN ID [0]: 100`
		Specifies the location of the new image. This command prompts you for the URL for the TFTP server, the management interface IP address and netmask, gateway address, and VLAN ID (ASA 5505 only). These network parameters are configured in ROMMON; the network parameters you configured in the module application configuration are not available to ROMMON, so you must set them separately here.
		
		You can view the recovery configuration using the **show module 1 recover** command.
		
		In multiple context mode, enter this command in the system execution space.

	Command	Purpose
Step 2	**hw-module module 1 recover boot** **Example:** `hostname# hw-module module 1 recover boot`	Transfers the image from the TFTP server to the module and restarts the module.
Step 3	**show module 1 details** **Example:** `hostname# show module 1 details`	Checks the progress of the image transfer and module restart process. The Status field in the output indicates the operational status of the module. A module operating normally shows a status of "Up." While the ASA transfers an application image to the module, the Status field in the output reads "Recover." When the ASA completes the image transfer and restarts the module, the newly transferred image is running.

Resetting the Password

You can reset the module password to the default. The default password is cisco. After resetting the password, you should change it to a unique value using the module application.

Resetting the module password causes the module to reboot. Services are not available while the module is rebooting.

If you cannot connect to ASDM with the new password, restart ASDM and try to log in again. If you defined a new password and still have an existing password in ASDM that is different from the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

To reset the module password to the default of cisco, perform the following steps.

Detailed Steps

Step 1 From the ASDM menu bar, choose **Tools > CSC Password Reset**.

The Password Reset confirmation dialog box appears.

Step 2 Click **OK** to reset the password to the default.

A dialog box displays the success or failure of the password reset.

Step 3 Click **Close** to close the dialog box.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Detailed Steps

Command	Purpose
hw-module module 1 reload **Example:** `hostname# hw-module module 1 reload`	Reloads the module software.
hw-module module 1 reset **Example:** `hostname# hw-module module 1 reset`	Performs a reset, then reloads the module.

Shutting Down the Module

If you restart the ASA, the module is not automatically restarted. To shut down the module, perform the following steps at the ASA CLI.

Detailed Steps

Command	Purpose
hw-module module 1 shutdown **Example:** `hostname# hw-module module 1 shutdown`	Shuts down the module.

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.	*Cisco Content Security and Control SSM Administrator Guide*
Accessing ASDM for the first time and assistance with the Startup Wizard.	*Cisco ASA 5500 Series Quick Start Guide*
Assistance with SSM hardware installation and connection to the ASA.	hardware guide
Accessing ASDM for the first time and assistance with the Startup Wizard.	*Cisco ASA 5500 Series Quick Start Guide*
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.	*Cisco Content Security and Control SSM Administrator Guide*
Technical Documentation, Marketing, and Support-related information.	See the following URL: http://www.cisco.com/en/US/products/ps6823/index.html.

Feature History for the CSC SSM

Table 83-2 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 83-2 Feature History for the CSC SSM

Feature Name	Platform Releases	Feature Information
CSC SSM	7.0(1)	The CSC SSM runs Content Security and Control software, which provides protection against viruses, spyware, spam, and other unwanted traffic. The CSC Setup Wizard enables you to configure the CSC SSM in ASDM. We introduced the following screen: Configuration > Trend Micro Content Security > CSC Setup.
CSC SSM	8.1(1) and 8.1(2)	This feature is not supported on the ASA 5580.

Table 83-2 Feature History for the CSC SSM (continued)

Feature Name	Platform Releases	Feature Information
CSC syslog format	8.3(1)	CSC syslog format is consistent with the ASA syslog format. Syslog message explanations have been added to the *Cisco Content Security and Control 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.
Clearing CSC events	8.4(1)	Support for clearing CSC events in the Latest CSC Security Events pane has been added. We modified the following screen: Home > Content Security.
CSC SSM	8.4(2)	Support for the following features has been added: • HTTPS traffic redirection: URL filtering and WRS queries for incoming HTTPS connections. • Configuring global approved whitelists for incoming and outgoing SMTP and POP3 e-mail. • E-mail notification for product license renewals. We modified the following screens: Configuration > Trend Micro Content Security > Mail > SMTP. Configuration > Trend Micro Content Security > Mail > POP3. Configuration > Trend Micro Content Security > Host/Notification Settings. Configuration > Trend Micro Content Security > CSC Setup > Host Configuration.

PART 15

Configuring VPN

VPN Wizards

The ASA provides Secure Socket Layer (SSL) remote access connectivity from almost any Internet-enabled location using only a Web browser and its native SSL encryption. Clientless, browser-based VPN lets users establish a secure, remote-access VPN tunnel to the adaptive 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 ASA for remote users with full VPN tunneling to corporate resources. Without a previously-installed client, remote users enter the IP address in their browser of an interface configured to accept clientless VPN connections. The ASA downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure connection and either remains or uninstalls itself (depending on the ASA configuration) when the connection terminates. In the case of a previously installed client, when the user authenticates, the ASA examines the revision of the client and upgrades the client as necessary.

With the addition of IKEv2 support in release 8.4, the end user can have the same experience independent of the tunneling protocol used by the AnyConnect client session. This addition allows other vendors' VPN clients to connect to the ASAs. This support enhances security and complies with the IPsec remote access requirements defined in federal and public sector mandates.

The VPN wizard lets you configure basic LAN-to-LAN and remote access VPN connections and assign either preshared keys or digital certificates for authentication. Use ASDM to edit and configure advanced features.

VPN Overview

The ASA 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.

For LAN-to-LAN connections using both IPv4 and IPv6 addressing, the security appliance supports VPN tunnels if both peers are Cisco ASA 5500 series security appliances, and if both inside networks have matching addressing schemes (both IPv4 or both IPv6). This is also true if both peer inside networks are IPv6 and the outside network is IPv6.

The secure connection is called a tunnel, and the ASA uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The ASA 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 four VPN wizards described in this section are as follows:

- IPsec IKEv1 Remote Access Wizard
- IPsec Site-to-Site VPN Wizard
- AnyConnect VPN Wizard
- Clientless SSL VPN Wizard

IPsec IKEv1 Remote Access Wizard

Use the IKEv1 Remote Access Wizard to configuresecure remote access for VPN clients, such as mobile users, and to identify the interface that connects to the remote IPsec peer.

Fields

- VPN Tunnel Interface—Choose the interface that establishes a secure tunnel with the remote IPsec peer. If the ASA 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.

Remote Access Client

Remote access users of various types can open VPN tunnels to this ASA. Choose the type of VPN client for this tunnel.

Fields

- VPN Client Type
  - Cisco VPN Client, Release 3.x or higher, or an Easy VPN Remote product.
  - Microsoft Windows client using L2TP over IPsec—Specify the PPP authentication protocol. The choices are PAP, CHAP, MS-CHAP-V1, MS-CHAP-V2, and EAP-PROXY:

    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 ASA to proxy the PPP authentication process to an external RADIUS authentication server.

    If a protocol is not specified on the remote client, do no specify it.

– Specify if the client will send tunnel group name as username@tunnelgroup.

VPN Client Authentication Method and Tunnel Group 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 ASA 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 an alphanumeric string between 1 and 128 characters.
  - Certificate—Click to use certificates for authentication between the local ASA 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 ASA.

    You can efficiently manage the security keys used to establish an IPsec tunnel with digital certificates. 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.
  - 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 connection policy that you configure with this VPN wizard specifies an authentication method and uses the ASA Default Group Policy.

Client Authentication

Use the Client Authentication pane to select the method by which the ASA authenticates remote users.

Fields

Select one of the following options:

- Authenticate using the local user database—Click to use authentication internal to the ASA. Use this method for environments with a small, stable number of users. The next pane lets you create accounts on the ASA 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.

User Accounts

Use the User Accounts pane to add new users to the ASA internal user database for authentication purposes.

Fields

- 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.
- Delete—To remove a user from the database, highlight the appropriate username and click **Delete**.

Address Pool

Use the Address Pool pane to configure a pool of local IP addresses that the ASA assigns to remote VPN clients.

Fields

- Tunnel Group Name—Displays the name of the connection profile (tunnel group) to which this address pool applies. You set this name in the VPN Client and Authentication Method pane (step 3).
- 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.

Attributes Pushed to Client (Optional)

Use the Attributes Pushed to Client (Optional) pane to have the ASA pass information about DNS and WINS servers and the default domain name to remote access clients.

Fields

- 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.

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 ASA uses this algorithm to derive the encryption and hash keys.

Fields

- Encryption—Select the symmetric encryption algorithm the ASA uses to establish the Phase 1 SA that protects Phase 2 negotiations. The ASA 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 ASA 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 ASA 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.

IPsec Settings (Optional)

Use the IPsec Settings (Optional) pane to identify local hosts/networks which do not require address translation. By default, the ASA 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

- Interface—Choose the name of the interface that connects to the hosts or networks you have selected.
- Exempt Networks—Select the IP address of the host or network that you want to exempt from the chosen interface network.
- 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 ASA 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 ASA. All other traffic travels unencrypted directly to the Internet without involving the ASA.
- 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).

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**.

IPsec Site-to-Site VPN Wizard

Use this wizard to set up new site-to-site VPN tunnels. A tunnel between two devices is called a site-to-site tunnel and is bidirectional. A site-to-site VPN tunnel protects the data using the IPsec protocol.

Peer Device Identification

Identify the peer VPN device by its IP address and the interface used to access the peer.

Fields

- Peer IP Address—Configure the IP address of the peer device.
- VPN Access Interface—Use the drop-down to specify the interface for the site-to-site tunnel.

IKE Version

ASA supports both version 1 and version 2 of the IKE (Internet Key Exchange) protocol. This step lets you decide which version or versions to support in this connection profile.

Fields

- IKEv1
- IKEv2

Traffic to Protects

This step lets you identify the local network and remote network These networks protect the traffic using IPsec encryption.

Fields

- Network Type—Choose IPv4 or IPv6.
- Local Networks—Identify the host used in the IPsec tunnel.
- Remote Networks—Identify the networks used in the IPsec tunnel.

Authentication Methods

This step lets you configure the methods to authenticate with the peer device.

Fields

IKE version 1

- Pre-shared Key—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.

- Device Certificate—Click to use certificates for authentication between the local ASA and the remote IPsec peer.

  You can efficiently manage the security keys used to establish an IPsec tunnel with digital certificates. 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.

  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.

IKE version 2

- Local Pre-shared Key—Specify IPsec IKEv2 authentication methods and encryption algorithms.
- Local Device Certificate—Authenticates VPN access through the security appliance.
- Remote Peer Pre-shared Key—Click to use a preshared key for authentication between the local ASA and the remote IPsec peer.
- Remote Peer Certificate Authentication—When checked, the peer device is allowed to use the certificate to authenticate itself to this device.

Encryption Algorithm

This step lets you select the types of encryption algorithms used to protect the data.

Fields

IKE version 1

- IKE Policy—Specify IKEv1 authentication methods.
- IPsec Proposal—Specify IPsec encryption algorithms.

IKE version 2

- IKE Policy—Specify IKEv2 authentication methods.
- IPsec Proposal—Specify IPsec encryption algorithms.

Miscellaneous

You can enable or disable Perfect Forward Secrecy (PFS). PFS ensures that the key for a given IPsec SA was not derived from any other secret. PFS makes it difficult to break a key by deriving from other keys.

Fields

- 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.
- Enable Perfect Forward Secrecy (PFS)—Ensures the key for a given IPsec SA was not derived from any other secret.
- 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).
- Exempt ASA side host/network from address translation—Use the drop-down to choose a host or network to be excluded from address translation.

Summary

Provides a summary of your selections from the previous wizard windows. The supported VPN protocols are included in the summary as well as the IKE version chosen on the VPN Connection Type window.

AnyConnect VPN Wizard

Use this wizard to configure ASA to accept VPN connections from the AnyConnect VPN client. This wizard configures either IPsec (IKEv2) or SSL VPN protocols for full network access. The ASA automatically uploads the AnyConnect VPN client to the end user's device when a VPN connection is established.

Warn the user that running the wizard does not mean the IKEv2 profile automatically applies in predeployment scenarios. Either provide a pointer or the steps necessary to successfully predeploy IKEv2.

Connection Profile Identification

The connection profile identification is used to identify the ASA to the remote acess users.

Fields

- Connection Profile Name—Provide a name that the remote access users will access for VPN connections.
- VPN Access Interface—Choose an interface that the remote access users will access for VPN connections.

VPN Protocols

Specify the VPN protocol allowed for this connection profile.

The AnyConnect client defaults to SSL. If you enable IPsec as a VPN tunnel protocol for the connection profile, you must also create and deploy a client profile with IPsec enabled using the profile editor from ASDM, and deploy the profile.

If you predeploy instead of weblaunch the AnyConnect client, the first client connection uses SSL, and receives the client profile from the ASA during the session. For subsequent connections, the client uses the protocol specified in the profile, either SSL or IPsec. If you predeploy the profile with IPsec specified with the client, the first client connection uses IPsec. For more information about predeploying a client profile with IPsec enabled, see the AnyConnect Secure Mobility Client Administrator Guide.

Fields

- SSL
- IPsec (IKE v2)
- Device Certificate—Identifies the ASA to the remote access clients.

Note Some AnyConnect features (such as always on, IPsec/IKEv2) require a valid device certificate on the ASA.

- Manage—Choosing **Manage** opens the Manage Identity Certificates window.
  - Add—Choose **Add** to add an identity certificate and its details.
  - Show Details—If you choose a particular certificate and click **Show Details**, the Certificate Details window appears and provides who the certificate was issued to and issued by, as well as specifics about its serial number, usage, associated trustpoints, valid timeframe, and so on.
  - Delete—Highlight the certificate you want to remove and click **Delete**.
  - Export—Highlight the certificate and click **Export** to export the certificate to a file with or without an encryption passphrase.
  - Enroll ASA SSL VPN with Entrust—Gets your Cisco ASA SSL VPN appliance up and running quickly with an SSL Advantage digitial certificate from Entrust.

Client Images

ASA can automatically upload the latest AnyConnect package to the client device when it accesses the enterprise network. You can use a regular expression to match the user agent of a browser to an image. You can also minimize connection setup time by moving the most commonly encountered operation system to the top of the list.

Fields

- Add
- Replace
- Delete

Authentication Methods

Specify authentication information on this screen.

Fields

- AAA server group—Enable to let the ASA contact a remote AAA server group to authenticate the user. Select a AAA server group from the list of pre-configured groups or click **New** to create a new group.
- Local User Database Details—Add new users to the local database stored on the ASA.
  - 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.

Client Address Assignment

Provide a range of IP addresses to remote SSL VPN users.

Fields

- IPv4 Address Pools—SSL VPN clients receive new IP addresses when they connect to the ASA. Clientless connections do not require new IP addresses. Address Pools define a range of addresses that remote clients can receive. Select an existing IP Address Pool or click **New** to create a new pool.

  If you select **New**, you will have to provide a starting and ending IP address and subnet mask.
- IPv6 Address Pool—Select an existing IP Address Pool or click **New** to create a new pool.

Note IPv6 address pools can not be created for IKEv2 connection profiles.

Network Name Resolution Servers

This step lets you specify which domain names are resolved for the remote user when accessing the internal network.

Fields

- DNS Servers—Enter the IP address of the DNS server.
- WINS Servers—Enter the IP address of the WINS server.
- Domain Name—Type the default domain name.

NAT Exempt

If network translation is enabled on the ASA, the VPN traffic must be exempt from this translation.

Fields

- Exempt VPN traffic from network address translation

AnyConnect Client Deployment

You can install the AnyConnect client program to a client device with one of the following two methods:

- Web launch—Installs automatically when accessing the ASA using a web browser.
- Pre-deployment—Manually installs the AnyConnect client package.

Fields

- Allow Web Launch—A global setting that affects all connections. If it is unchecked (disallowed), AnyConnect SSL connections and clientless SSL connections do not work.

For pre-deployment, the disk0:/test2_client_profile.xml profile bundle contains an .msi file, and you must include this client profile from the ASA in your AnyConnect package to ensure IPsec connection functions as expected.

Summary

Provides a summary of your selections from the previous wizard windows. The supported VPN protocols are part of the summary as well as the IKE version chosen.

Clientless SSL VPN Wizard

This wizard enables clientless, browser-based connections for specific, supported internal resources through a portal page.

SSL VPN Interface

Provide a connection profile and the interface that SSL VPN users connect to.

Fields

- Connection Profile Name
- SSL VPN Interface—The interface users access for SSL VPN connections.
- Digital Certificate—Specifies what the security appliance sends to the remote web browser to authenticate the ASA.
  - Certificate—Choose from the drop-down menu.
- Accessing the Connection Profile
  - Connection Group Alias/URL—The group alias is chosen during login from the Group drop-down list. This URL is entered into the web browser.
  - Display Group Alias list at the login page

User Authentication

Specify authentication information on this screen.

Fields

- Authenticate using a AAA server group—Enable to let the ASA 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 ASA.
  - 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.

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—Enables you to create a new group policy. Provide a name for the new policy.
- Modify existing group policy—Select an existing group policy to modify.

Bookmark List

Configure a list of group intranet websites that appear in the portal page as links. Some examples include https://intranet.acme.com, rdp://10.120.1.2, vnc://100.1.1.1 and so on.

Fields

- Bookmark List
- Manage

Summary

Provides a summary of your selections from the previous wizard windows.

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 ASA 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 ASAs 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, page 86-1
- Creating IKE Policies, page 86-5
- Configuring IPsec, page 86-10
- Configuring Load Balancing, page 86-20
- Setting Global NAC Parameters, page 86-29
- Configuring Network Admission Control Policies, page 86-30

Setting IKE Parameters

Configuration > VPN > IKE > Global 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 ASA 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 ASA 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 ASA.
- 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 ASA feature only. It does not work for LAN-to-LAN connections.

- The ASA 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 ASA 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 ASA 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 ASA. The client configuration must include at least one of the ports you set for the ASA.

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 ASA shutdown or reboot, session idle timeout, maximum connection time exceeded, or administrator cut-off.

The ASA 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 ASA 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 ASA reboot to occur only when all active sessions have terminated voluntarily. This feature is disabled by default.

Preventing DoS Attacks by Limiting IKEv2 Open SAs

You can prevent denial-of-service (DoS) attacks for IPsec IKEv2 connections by always cookie challenging incoming SAs or by limiting the number of open SAs and cookie challenge any additional connections, or by By default, the ASA does not limit the number of open SAs and never cookie challenges SAs. You can also limit the number of SAs allowed, which stops further connections from negotiating to protect against memory and/or CPU attacks that the cookie-challenge feature may be unable to thwart and protects the current connections.

With a DoS attack, an attacker initiates the attack when the peer device sends an SA initiate packet and the ASA sends its response, but the peer device does not respond further. If the peer device does this continually, all the allowed SA requests on the ASA can be used up until it stops responding.

Enabling a threshold percentage for cookie challenging limits the number of open SA negotiations. For example, with the default setting of 50%, when 50% of the allowed SAs are in-negotiation (open), the ASA cookie challenges any additional SA initiate packets that arrive. For the Cisco ASA 5580 with 10000 allowed IKEv2 SAs, after 5000 SAs become open, any more incoming SAs are cookie-challenged.

If used in conjunction with the ***Number of SAs in Negotiation***, or the *Maximum Number of SAs Allowed*, configure the cookie-challenge threshold lower than these settings for an effective cross-check.

Fields

- Enable IKE—Shows IKE status for all configured interfaces.
  - Interface—Displays names of all configured ASA 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 ASA 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 ASA 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 ASA postpone a scheduled reboot until all active sessions terminate.
- IKEv2 Specific Settings—These settings apply only to IPsec IKEv2 connections and limit the number of open SAs. By default, the ASA does not limit the number of open SAs:
  - Cookie Challenge—Enables the ASA to send cookie challenges to peer devices in response to SA initiate packets.

  - % threshold before incoming SAs are cookie challenged—The percentage of the total allowed SAs for the ASA that are in-negotiation, which triggers cookie challenges for any future SA negotiations. The range is zero to 100%. The default is 50%.
  - Number of Allowed SAs in Negotiation—Limits the maximum number of SAs that can be in negotiation at any time. If used in conjunction with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check.
  - Maximum Number of SAs Allowed—Limits the number of allowed IKEv2 connections on the ASA. By default, the limit is the maximum number of connections specified by the license.

Creating IKE Policies

Configuration > VPN > IKE > Policies

Each IKE negotiation is divided into two sections called Phase1 and Phase 2.

Phase 1 creates the first tunnel, which protects later IKE negotiation messages. Phase 2 creates the tunnel that protects data.

To set the terms of the IKE negotiations, you create one or more IKE policies, which include the following:

- A unique priority (1 through 65,543, with 1 the highest priority).
- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.
- An HMAC method to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to establish the strength of the of the encryption-key-determination algorithm. The ASA uses this algorithm to derive the encryption and hash keys.
- A limit for how long the ASA uses an encryption key before replacing it.

For IKEv1, you can only enable one setting for each parameter. For IKEv2, each proposal can have multiples settings for Encryption, D-H Group, Integrity Hash, and PRF Hash.

If you do not configure any IKE policies, the ASA uses the default policy, which is always set to the lowest priority, and which contains the 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

- IKEv1 Policies—Displays parameter settings for each configured IKE policy.
  - Priority #—Shows the priority of the policy.
  - Encryption—Shows the encryption method.
  - Hash—Shows the hash 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 IKEv1 policy.
- IKEv2 Policies—Displays parameter settings for each configured IKEv2 policy.
  - Priority #—Shows the priority of the policy.
  - Encryption—Shows the encryption method.
  - Integrity Hash—Shows the hash algorithm.
  - PRF Hash—Shows the pseudo random function (PRF) hash algorithm.
  - D-H Group—Shows the Diffie-Hellman group.
  - Lifetime (secs)—Shows the SA lifetime in seconds.
- Add/Edit/Delete—Click to add, edit, or delete an IKEv2 policy.

Add/Edit IKEv1 Policy

Configuration > VPN > IKE > Policies > Add/Edit IKEv1 Policy

Fields

Priority #—Type a number to set a priority for the IKE policy. The range is 1 to 65535, 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 ASA 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 1or 5.
2	Group 2 (1024-bit)	
5	Group 5 (1536-bit)	

Lifetime (secs)—Either check *Unlimited* or enter an integer for the SA lifetime. The default is 86,400 seconds or 24 hours. With longer lifetimes, the ASA sets up future IPsec security associations less 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 ASA accepts the following values:.

120 - 86,400 seconds

2 - 1440 minutes

1 - 24 hours

1 day

Add/Edit IKEv2 Policy (Proposal)

Configuration > VPN > IKE > Policies > Add/Edit IKEv2 Policy

Fields

Priority #—Type a number to set a priority for the IKEv2 policy. The range is 1 to 65535, 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	Specifies 56-bit DES-CBC encryption for ESP.
3des	(Default) Specifies the triple DES encryption algorithm for ESP.
aes	Specifies AES with a 128-bit key encryption for ESP.
aes-192	Specifies AES with a 192-bit key encryption for ESP.
aes-256	Specifies AES with a 256-bit key encryption for ESP.
aes-gcm	Specifies AES-GCM/GMAC 128-bit support for symmetric encryption and integrity.
aes-gcm-192	Specifies AES-GCM/GMAC 192-bit support for symmetric encryption and integrity.
aes-gcm-256	Specifies AES-GCM/GMAC 256-bit support for symmetric encryption and integrity.
NULL	Indicates no encryption.

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)	
14	Group 14	
19	Group 19	
20	Group 20	
21	Group 21	
24	Group 24	

Integrity Hash—Choose the hash algorithm that ensures data integrity for the ESP protocol. It ensures that a packet comes from whom you think it comes from, and that it has not been modified in transit.

sha md5	SHA 1 MD5	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.
sha256	SHA 2, 256-bit digest	Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
sha384	**SHA 2, 384-bit digest**	Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
sha512	**SHA 2, 512-bit digest**	Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.
null		Indicates that AES-GCM or AES-GMAC is configured as the encryption algorithm. You must choose the null integrity algorithm if AES-GCM has been configured as the encryption algorithm.

Pseudo-Random Function (PRF)—Specify the PRF used for the construction of keying material for all of the cryptographic algorithms used in the SA..

sha md5	SHA-1 MD5	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.
sha256	SHA 2, 256-bit digest	Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
sha384	**SHA 2, 384-bit digest**	Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
sha512	**SHA 2, 512-bit digest**	Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.

Lifetime (secs)—Either check *Unlimited* or enter an integer for the SA lifetime. The default is 86,400 seconds or 24 hours. With longer lifetimes, the ASA 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 ASA accepts the following values:.

120 - 86,400 seconds

2 - 1440 minutes

1 - 24 hours

1 day

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 ASA 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 ASA 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 ASA 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 ASA 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.

Configuring IPsec

The ASA 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 (IPv4 or IPv6), 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 ASA can function as initiator or responder. In IPsec client-to-LAN connections, the ASA 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 ASA 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

Configuration > VPN > IPSec > IPSec Rules

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 ASA 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 ASA maps the inside host's address to an address from the pool. After a host creates an outbound connection, the ASA 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.

Creating an IPsec Rule/Tunnel Policy (Crypto Map) - Basic Tab

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - Basic Tab

Use this pane to define a new Tunnel Policy for an IPsec rule. The values you define here appear in the IPsec Rules table after you click OK. All rules are enabled by default as soon as they appear in the IPsec Rules table.

The Tunnel Policy 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.
- IKE Proposals (Transform Sets)--Specifies IKEv1 and IKEv2 IPsec proposals:
  - IKEv1 IPsec Proposal—Choose the proposal (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 proposals in the list box. You can add a maximum of 11 proposals to a crypto map entry or a dynamic crypto map entry.
  - IKEv2 IPsec Proposal—Choose the proposal (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 proposals in the list box. You can add a maximum of 11 proposals 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. If you choose Originate Only, you can specify up to 10 redundant peers. For uni-directional, you can specify originate only or answer only, and neither are enabled by default.
  - 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 ASA 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.
  - Group 14= Use perfect forward secrecy and use Diffie-Hellman Group 14 for IKEv2.
  - Group 19= Use perfect forward secrecy and use Diffie-Hellman Group 19 for IKEv2 to support ECDH.
  - Group 20= Use perfect forward secrecy and use Diffie-Hellman Group 20 for IKEv2 to support ECDH.
  - Group 21= Use perfect forward secrecy and use Diffie-Hellman Group 21 for IKEv2 to support ECDH.
  - Group 24= Use perfect forward secrecy and use Diffie-Hellman Group 24 for IKEv2.

Creating IPsec Rule/Tunnel Policy (Crypto Map) - Advanced Tab

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - Advanced Tab

Fields

- 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 dynamic routing protocols such as Open Shortest Path First (OSPF), or Enhanced Interior Gateway Routing Protocol (EIGRP), if you run ASA, or Routing Information Protocol (RIP) for remote VPN Clients or LAN to LAN sessions.
- Security Association Lifetime Settings—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 Type Only Settings—Specifies parameters for static tunnel policies.
  - Device Certificate—Choose the certificate to use. If you choose something other than None (Use Preshared Keys), which is the default. The Send CA certificate chain check box becomes active when you select something other than None.
  - Send CA certificate chain—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).
- ESP v3—Specify whether incoming ICMP error messages are validated for cryptography and dynamic cryptography maps, set the per-security association policy, or enable traffic flow packets:
  - Validate incoming ICMP error messages—Choose whether to validate those ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.
  - Enable Do Not Fragment (DF) policy—Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header. Choose one of the following:

    Clear DF bit—Ignores the DF bit.

    Copy DF bit—Maintains the DF bit.

    Set DF bit—Sets and uses the DF bit.
  - Enable Traffic Flow Confidentiality (TFC) packets—Enable dummy TFC packets that mask the traffic profile which traverses the tunnel.

Note You must have an IKE v2 IPsec proposal set on the Tunnel Policy (Crypto Map) Basic tab before enabling TFC.

Use the Burst, Payload Size, and Timeout parameters to generate random length packets at random intervals across the specified SA.

Creating IPsec Rule/Traffic Selection Tab

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - 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.

Pre-Fragmentation

Configuration > VPN > IPSec > 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 ASA and the client rejects or drops IP fragments. For example, suppose a client wants to FTP get from an FTP server behind a ASA. The FTP server transmits packets that when encapsulated would exceed the ASA's MTU size on the public interface. The selected options determine how the ASA processes these packets. The pre-fragmentation policy applies to all traffic travelling out the ASA public interface.

The ASA encapsulates all tunneled packets. After encapsulation, the ASA 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 ASA fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the ASA 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 ASA 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.

Edit IPsec Pre-Fragmentation Policy

Configuration > VPN > IPSec > Pre-Fragmentation > 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 ASA fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the ASA 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.

IPsec Transform Sets

Configuration > VPN > 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

- **IKEv1 IPsec Proposals (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.
- **IKEv2 IPsec Proposals**—Shows the configured transform sets.
  - **Name**—Shows the name of the **IKEv2 IPsec Proposal**.
  - **Encryption**—Shows the Encapsulating Security Protocol (ESP) encryption algorithms for the **IKEv2 IPsec Proposal**. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
  - **Integrity Hash**—Shows the hash algorithm that ensures data integrity for the ESP protocol. It ensures that a packet comes from whom you would expect and that no modifications were made in transit. It ensures that a packet comes from who you would expect and that no modifications were made in transit. You must choose the null integrity algorithm if AES-GCM/GMAC has been configured as the encryption algorithm.
- **Add**—Opens the Add IPsec Proposal dialog box, in which you can add a new proposal.
- **Edit**—Opens the Edit IPsec Proposal dialog box, in which you can modify an existing proposal.
- **Delete**—Removes the selected proposal. There is no confirmation or undo.

Add/Edit IPsec Proposal (Transform Set)

(Configuration > VPN > IPSec > Transform Sets > Add/Edit IPsec_Proposal_(Transform Set)

Use this pane to add or modify an IPsec IKEv1 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."

Add/Edit IPsec Proposal

Configuration > VPN > IPSec > Transform Sets > Add/Edit IPsec_Proposal

Use this pane to add or modify an IPsec IKEv2 proposal. A proposal is a set of operations done on a data flow to provide data authentication, data confidentiality, and data compression. For example, one proposal is the ESP protocol with 3DES encryption and the HMAC-MD5 authentication algorithm (ESP-3DES-MD5).

Fields

- **Name**—Specifies a name for this proposal.
- **Encryption**—Choose the Encapsulating Security Protocol (ESP) encryption algorithms for the proposal. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
- **Integrity Hash**—Choose the ESP authentication algorithms for the proposal. The hash algorithm ensures data integrity for the ESP protocol. It ensures that a packet comes from whom you think it comes from, and that it has not been modified in transit.

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."

Configuring Load Balancing

If you have a remote-client configuration in which you are using two or more ASAs 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.

The following sections explain load balancing:

- Creating Virtual Clusters
- Geographical Load Balancing
- Mixed Cluster Scenarios
- Comparing Load Balancing to Failover
- Load Balancing Licensing Requirements
- Load Balancing Prerequisites
- Eligible Clients

- Configuring VPN Cluster Load Balancing with the High Availability and Scalability Wizard
- Configuring Load Balancing (Without the Wizard)

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 connection requests 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 ASA 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 ASAs 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 ASAs 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 ASAs or same release ASAs and VPN 3000 concentrators or both can support only IPsec sessions. In such a configuration, however, the ASAs 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 ASA 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.

Geographical Load Balancing

In a load balancing environment where the DNS resolutions are being changed at regular intervals, you must carefully consider how to set the time to live (TTL) value. For the DNS load balance configuration to work successfully with AnyConnect, the ASA name to address mapping must remain the same from the time the ASA is selected until the tunnel is fully established. If too much time passes before the credentials are entered, the lookup restarts and a different IP address may become the resolved address. If the DNS mapping changes to a different ASA before the credentials are entered, the VPN tunnel fails.

Geographical load balancing for VPN often uses a Cisco Global Site Selector (GSS). The GSS uses DNS for the load balancing, and the time to live (TTL) value for DNS resolution is defaulted to 20 seconds. You can significantly decrease the likelihood of connection failures if you increase the TTL value on the GSS. Increasing to a much higher value allows ample time for the authentication phase when the user is entering credentials and establishing the tunnel.

To increase the time for entering credentials, you may also consider disabling Connect on Start Up.

Note Using a Cisco System VPN client, the geographical load balancing is not impacted, and the 20 second default setting is acceptable.

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 ASA 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 ASAs 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 ASAs and VPN 3000 Concentrators. Some of the ASA 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 ASA 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 ASA 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) ASA. The following scenario illustrates this dilemma.

This scenario is similar to the previous one, in that the cluster consists of a mixture of ASAs and VPN 3000 Concentrators. Some of the ASA 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 ASAs 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 ASA 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 ASAto 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 Licensing Requirements

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.

Eligible Clients

Load balancing is effective only on remote sessions initiated with the following clients:

- Cisco AnyConnect Secure Mobility 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 ASA on which load balancing is enabled, but they cannot participate in load balancing.

Load Balancing Prerequisites

- You must have first configured the ASA's public and private interfaces before configuring load balancing. To do so select **Configuration > Device Setup > Interfaces**. See Chapter 15, "Starting Interface Configuration (ASA 5510 and Higher)" or Chapter 16, "Starting Interface Configuration (ASA 5505)" for more information.
- You must have previously configured 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.

Configuring VPN Cluster Load Balancing with the High Availability and Scalability Wizard

If you have a remote-client configuration in which you are using two or more ASAs 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, which directs session traffic to the least loaded device, thereby distributing the load among all devices. Load balancing makes efficient use of system resources and provides increased performance and system availability.

Use the VPN Cluster Load Balancing Configuration screen to set required parameters for a device to participate in a load balancing cluster.

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 each 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 configured as an Easy VPN client. All other clients, including LAN-to-LAN connections, can connect to a ASA on which load balancing is enabled, but these clients cannot participate in load balancing.

To implement load balancing, you logically group together two or more devices on the same private LAN-to-LAN network into a virtual cluster by performing the following steps:

Step 1 Choose **Wizards > High Availability and Scalability**.

Step 2 In the Configuration Type screen, click **Configure VPN Cluster Load Balancing**, and click **Next**.

Step 3 Choose the single IP address that represents the entire virtual cluster. Specify an IP address that is within the public subnet address range shared by all the ASAs in the virtual cluster.

Step 4 Specify 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 that you want to use for load balancing.

Step 5 To enable IPsec encryption and ensure that all load-balancing information communicated between the devices is encrypted, check the **Enable IPsec Encryption** check box. You must also specify and verify a shared secret. The ASAs in the virtual cluster communicate via LAN-to-LAN tunnels using IPsec. To disable IPsec encryption, uncheck the **Enable IPsec Encryption** 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, an error message appears 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, an error message appears when you check the **Participate in Load Balancing Cluster** check box, and encryption is not enabled for the cluster.

Step 6 Specify the shared secret to between IPsec peers when you enable IPsec encryption. The value that you enter appears as consecutive asterisk characters.

Step 7 Specify 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 startup or when an existing master fails. The higher the priority set (for example, 10), the more likely that this device will become 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 assumes 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.

Step 8 Specify the name or IP address of the public interface for this device.

Step 9 Specify the name or IP address of the private interface for this device.

Step 10 Check the **Send FQDN to client instead of an IP address when redirecting** check box to have the VPN cluster master 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.

Step 11 Click **Next**. Review your configuration in the Summary screen.

Step 12 Click **Finish**.

The VPN cluster load balancing configuration is sent to the ASA.

Configuring Load Balancing (Without the Wizard)

The pane Load Balancing pane (Configuration > Remote Access VPN > Load Balancing) lets you enable load balancing on the ASA. 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.

Step 1 Select **Configuration > Remote Access VPN > Load Balancing.**

Step 2 Check **Participate in Load Balancing** to indicate that this ASA is a participant in the load-balancing cluster

You must enable load balancing in this way on every ASA participating in load balancing.

Step 3 Configure the following fields in the **VPN Cluster Configuration** area. These values must be the same for the entire virtual cluster. All servers in the cluster must have an identical cluster configuration.

- **Cluster IPv4 Address**—Specifies the single IPv4 address that represents the entire IPv4 virtual cluster. Choose an IP address that is within the public subnet address range shared by all the ASAs 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.
- **Cluster IPv6 Address**—Specifies the single IPv6 address that represents the entire IPv6 virtual cluster. Choose an IP address that is within the public subnet address range shared by all the ASAs in the virtual cluster.

Note In the Cluster IPv4 Address and Cluster IPv6 Address fields, you can also specify the fully qualified domain name of the virtual cluster, provided that you have a DNS server group configured with at least one DNS server, and DNS lookup is enabled on one of the ASA's interfaces.

- **Enable IPsec Encryption**—Enables or disables IPsec encryption. If you check this box, you must also specify and verify a shared secret.The ASAs 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 between IPsec peers when you have enabled IPsec encryption. The value you enter in the box appears as consecutive asterisk characters.
- **Verify Secret**—Re-enter the shared secret. Confirms the shared secret value entered in the IPsec Shared Secret box.

Step 4 Configure the fields in the **VPN Server Configuration** area for a specific ASA:

- **Public Interface**—Specifies the name or IP address of the public interface for this device.
- **Private Interface**—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 IPv4 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.
- **NAT Assigned IPv6 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 ASA can send a fully qualified domain name (FQDN), using reverse DNS lookup, of a cluster device (another ASA 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.

  Note When using IPv6 and sending FQDNS down to client, those names must be resolvable by the ASA via DNS.

Enable Clientless SSL VPN Load Balancing Using FQDNs

Step 1 Enable the use of FQDNs for Load Balancing by checking the **Send FQDN to client instead of an IP address when redirecting** checkbox.

Step 2 Add an entry for each of your ASA outside interfaces into your DNS server, if such entries are not already present. Each ASA outside IP address should have a DNS entry associated with it for lookups. These DNS entries must also be enabled for Reverse Lookup.

Step 3 Enable DNS lookups on your ASA 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 ASA. To do this, click Add on this dialog box. This opens the Add DNS Server Group dialog box. Enter the IPv4 or IPv6 address of the DNS server you want to add; for example, 192.168.1.1 or 2001:DB8:2000::1.

Step 5 Click OK and Apply.

Setting Global NAC Parameters

The ASA 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 ASA.

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 ASA 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 ASA 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 ASA 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 ASA 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 ASA 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 ASA also starts this timer after it receives an Access Reject message from the ACS server. When the timer expires, the ASA 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 ASA 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 ASA 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.

Configuring Network Admission Control Policies

The NAC Policies table displays the Network Admission Control (NAC) policies configured on the ASA.

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 ASA triggers posture validation.

You can configure the ASA 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 ASA for application to the traffic on the tunnel.

In a *NAC Framework* configuration involving the ASA, 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 ASA 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 ASA, the ASA 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 ASA, 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 ASA 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 ASA 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 ASA 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 ASA starts this timer after each successful posture validation. The expiration of this timer triggers the next unconditional posture validation. The ASA 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 ASA 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 ASA 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 ASA 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 ASA 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.

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 ASA 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.

General VPN Setup

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, page 87-1
- Default Tunnel Gateway, page 87-4
- Group Policies, page 87-4
- Access Control List Manager, page 87-31
- Configuring AnyConnect VPN Client Connections, page 87-42
- Configuring AnyConnect VPN Connections, page 87-52
- Configuring AnyConnect Secure Mobility, page 87-64
- IPsec Remote Access Connection Profiles, page 87-73
- Add or Edit an IPsec Remote Access Connection Profile, page 87-74
- Mapping Certificates to IPsec or SSL VPN Connection Profiles, page 87-75
- System Options, page 87-102
- Zone Labs Integrity Server, page 87-103
- Easy VPN Remote, page 87-105
- Advanced Easy VPN Properties, page 87-107
- AnyConnect Essentials, page 87-108
- Configuring AnyConnect Host Scan, page 87-113
- Configuring Maximum VPN Sessions, page 87-119
- Configuring the Pool of Cryptographic Cores, page 87-119

Client Software

Configuration > VPN > General > Client Update

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 ASA 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.

Edit Client Update Entry

Configuration > VPN > General > Client Update > Edit Client Update Entry

The Edit Client Update dialog box lets you change information about VPN client revisions and URLs for the indicated client types. The clients must be running one of the revisions specified for the indicated client type. If not, the clients are notified that an upgrade is required.

Fields

- Client Type—(*Display-only*) Displays the client type selected for editing.
- VPN Client Revisions—Lets you type a comma-separated list of software or firmware images appropriate for this client. If the user 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.

Default Tunnel Gateway

Configuration > VPN > General > Default Tunnel Gateway

To configure the default tunnel gateway, click the **Static Route** link. The Configuration > Routing > Routing > Static Route dialog box opens.

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, 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.

You can configure these types of group policies:

- Configuring External Group Policies—An external group policy points the ASA to the RADIUS or LDAP server to retrieve much of the policy information that would otherwise be configured in an internal group policy. External group policies are configured the same way for Network (Client) Access VPN connections, Clientless SSL VPN connections, and Site-to-Site VPN connections.
- Configuring Network (Client) Access Internal Group Policies—These connections are initiated by a VPN client installed on the endpoint. The AnyConnect Secure Mobility Client and Cisco VPN IPsec client are examples of VPN clients. After the VPN client is authenticated, remote users can access corporate networks or applications as if they were on-site. The data traffic between remote users and the corporate network is secured by being encrypted when going through the Internet.
- Configuring Clientless SSL VPN Internal Group Policies—This is also known as browser-based VPN access. On successful login to the ASA's portal page, remote users can access corporate networks and applications from the links shown in the web pages. The data traffic between remote users and the corporate network is secured by traveling through SSL tunnel.
- Configuring Site-to-Site Internal Group Policies

Group Policy Pane Fields

Lists the currently configured group policies and Add, Edit, and Delete buttons to help you manage VPN group policies.

- 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.
- Assign—Lets you assign a group policy to one ore more connection profiles.
- 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.
- Connection Profiles/Users Assigned to—Lists the connection profiles and users configured directly on the ASA that are associated with this group policy.

Configuring External Group Policies

An external group policy points the ASA to the RADIUS or LDAP server to retrieve much of the policy information that would otherwise be configured in an internal group policy. External group policies are configured the same way for Network (Client) Access VPN connections, Clientless SSL VPN connections, and Site-to-Site VPN connections.

External group policies take their attribute values from the external server that you specify. For an external group policy, you must identify the RADIUS or LDAP server group that the ASA can query for attributes and specify the password to use when retrieving attributes from that external server group. If you are using an external authentication server, and if your external group-policy attributes exist in the same RADIUS server as the users that you plan to authenticate, you have to make sure that there is no name duplication between them.

Note External group names on the ASA refer to user names on the RADIUS server. In other words, if you configure external group X on the ASA, the RADIUS server sees the query as an authentication request for user X. So external groups are really just user accounts on the RADIUS server that have special meaning to the ASA. If your external group attributes exist in the same RADIUS server as the users that you plan to authenticate, there must be no name duplication between them.

The ASA supports user authorization on an external LDAP or RADIUS server. Before you configure the ASA to use an external server, you must configure the server with the correct ASA authorization attributes and, from a subset of these attributes, assign specific permissions to individual users. Follow the instructions in Appendix B, "Configuring an External Server for Authorization and Authentication" to configure your external server.

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.
- 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.
- Password—Specifies the password for this server group policy.

Adding an LDAP or RADIUS Server to a Network (Client) Access External Group Policy

Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > Add or Edit External Group Policy > New > RADIUS Server Group/New LDAP Server Group > Add AAA Server Group

The Add AAA Server Group dialog box lets you configure a new AAA server group. The Accounting Mode attribute applies only to RADIUS and TACACS+ protocols.

Fields

- Server Group—Specifies the name of the server group.
- Protocol—(*Display only*) Indicates whether this is a RADIUS or an LDAP server group.
- Accounting Mode—Indicates whether to use simultaneous or single accounting mode. In single mode, the ASA sends accounting data to only one server. In simultaneous mode, the ASA 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.

Configuring Network (Client) Access Internal Group Policies

Configure Network (Client) Access internal group policies for VPN connections made from AnyConnect Secure Mobility Clients or legacy Cisco IPsec VPN clients installed on an endpoint.

Configuring General Attributes for an 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 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.

You can configure the general attributes of an internal group policy by starting ASDM and selecting **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add or Edit Internal Group Policy > General.**

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.

  The IPsec VPN client supports full HTML for the banner. However, the clientless portal and the AnyConnect client support partial HTML. To ensure the banner displays properly to remote users, follow these guidelines:

  - For IPsec client users, use the /n tag.
  - For AnyConnect client users, use the <BR> tag.
  - For clientless users, use the <BR> tag.
- Address Pools—Specifies the name of one or more IPv4 address pools to use for this group policy. If the Inherit check box is checked, the group policy will use the IPv4 address pool specified in the Default Group Policy. See Configuring Local IP Address Pools, page 88-3 for information on adding or editing an IPv4 address pool.

  Select—Uncheck the Inherit checkbox to activate the Select command button. Click Select to open the 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. See Configuring Local IP Address Pools, page 88-3 for information on adding or editing an IPv6 address pool.

Note You can specify both an IPv4 and an IPv6 address pool for an internal group policy.

- More Options—Click the down arrows at the right of the field to display 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 an ASA; 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. If you are using the AnyConnect client, you must choose this protocol for Mobile User Security (MUS) to be supported.
  - IPsec IKEv1—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 Cisco VPN client-to-LAN connections can use IPsec IKEv1.
  - IPsec IKEv2—Supported by the AnyConnect Secure Mobility Client. AnyConnect connections using IPsec with IKEv2 provide advanced features such as software updates, client profiles, GUI localization (translation) and customization, Cisco Secure Desktop, and SCEP proxy.
  - 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.
- Filter—Specifies which unified access control list to use for an IPv4 or 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 ASA, based on criteria such as source address, destination address, and protocol. To configure filters and rules, see the ACL Manager 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.
- 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. See Defining Time Ranges, page 87-29 for more information.

- 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 ASA forwards all traffic from 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 ASA.

Note This feature works for HTTP connections, but not for FTP and CIFS.

- Connection Profile (Tunnel Group) Lock—This parameter permits remote VPN access only with the selected connection profile (tunnel group), and prevents access with a different connection profile. The default inherited value is None.
- 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, should we be so lucky). 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.

Configuring Server Attributes for an Internal Group Policy

Configure DNS servers, WINS servers and DHCP Scope in the Group Policy > Servers window.

Configuring a DNS Server for an Internal Group Policy

Use this procedure to configure a specific DNS server for a group policy.

Note This setting overrides the DNS setting configured on the ASDM in the **Configuration > Remote Access VPN > DNS** window.

Step 1 Select **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > Servers**.

Step 2 Unless you are editing the DefaultGroupPolicy, uncheck the DNS Servers **Inherit** checkbox.

Step 3 In the DNS Servers field, add the IPv4 or IPv6 addresses of the DNS servers you want this group to use.

You can specify up to four DNS server addresses: Up to two IPv4 addresses and two IPv6 addresses. If you specify more than one DNS server, the remote access client will attempt to use the DNS servers in the order you specify them in the field.

Step 4 Expand the **More Options** area by clicking the double down arrow in the More Options bar.

Step 5 If there is no default domain specified in the **Configuration > Remote Access VPN > DNS** window, you must specify the default domain in the **Default Domain** field. Use the domain name and top level domain for example, **example.com.**

Step 6 Click **OK**.

Step 7 Click **Apply**.

Configuring WINS Servers for an Internal Group Policy

Use this procedure to configure primary and secondary WINS servers. The default value in each case is none.

Step 1 Select **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > Servers**.

Step 2 Uncheck the WINS Servers **Inherit** checkbox.

Step 3 In the WINS Servers field, enter the IP addresses of the primary and secondary WINS servers. The first IP address you specify is that of the primary WINS server. The second (optional) IP address you specify is that of the secondary WINS server.

Step 4 Click OK.

Configuring Split Tunneling for an Internal Group Policy

Split tunneling enables you to route some network traffic from the endpoint through the VPN tunnel (encrypted) to the ASA and to route other network traffic from the endpoint outside the VPN tunnel (unencrypted or “in the clear”).

You implement split tunneling by creating a unified access control list and then requiring the list to be included in the VPN tunnel or requiring the list to be excluded from the VPN tunnel. You can specify both IPv4 and IPv6 addresses in the unified access control list.

The AnyConnect client and the legacy Cisco VPN client (the IPsec/IKEv1 client) behave differently when passing traffic to sites within the same subnet as the IP address assigned to the client by the ASA. The AnyConnect client passes traffic to all sites specified in the split tunneling policy you configured and to all sites that fall within the same subnet as the IP address assigned by the ASA. For example, if the IP address assigned to the client by the ASA is 10.1.1.1 with a mask of 255.0.0.0, the endpoint device passes all traffic destined to 10.0.0.0/8, regardless of the split tunneling policy. The legacy Cisco VPN client only passes traffic to addresses specified by the split-tunneling policy, regardless of the subnet assigned to the client by the ASA.

Therefore, in the unified access control list defined in the split tunnel IPv4 or IPv6 policy, use a netmask for the assigned IP address that properly references the expected local subnet.

Note In the following procedure, in all cases where there is an Inherit checkbox next to a field, leaving the Inherit check box checked means that the group policy you are configuring will use the same values for that field as the default group policy. Unchecking Inherit lets you specify new values specific to your group policy.

Step 1 Connect to the ASA using ASDM and select **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**.

Step 2 Click **Add** to add a new group policy or select an existing group policy and click **Edit**.

Step 3 Select **Advanced > Split Tunneling**.

Step 4 In the **DNS Names** field, enter the domain names that are to be resolved by AnyConnect via the tunnel. These names correspond to hosts in the private network. If split-include tunneling is configured, the network list must include the specified DNS servers. You can enter a full qualified domain name, IPv4 or IPv6 address in the field.

Step 5 The **Send All DNS Lookups Through Tunnel** field instructs the AnyConnect client to resolve all DNS addresses through the VPN tunnel (SSL or IPsec/IKEv2). Setting this option ensures that DNS traffic is not leaked to the physical adapter and disallows traffic in the clear. If DNS resolution fails, the address remains unresolved and the AnyConnect client does not try to resolve the address through public DNS servers. If you choose No (the default), the client sends DNS queries over the tunnel according to the split tunnel policy.

Step 6 Configure the **Policy** and **IPv6 Policy** fields—The Policy field defines the split tunneling policy for IPv4 network traffic. The IPv6 Policy field selects the split tunneling policy for IPv6 network traffic. Other than that difference, these fields have the same purpose.

You configure split-tunneling by unchecking the Inherit check box and choosing a split-tunneling policy. If you do not uncheck Inherit, your group policy uses the split tunneling settings defined in the default group policy, "DfltGrpPolicy." The default split tunneling policy setting in the default group policy is to Tunnel All Networks.

After unchecking Inherit, you can choose one of these policy options:

- **Exclude Network List Below**—This policy defines a list of networks to which traffic is sent in the clear. This feature is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel. This option applies only to the Cisco VPN client.
- **Tunnel Network List Below**—This policy tunnels all traffic from or to the networks specified in the Network List. This option enables split tunneling. It lets you create a network list of addresses to tunnel. Data to all other addresses travels in the clear and is routed by the remote user's Internet service provider.
- **Tunnel All Networks**—This policy specifies that no traffic goes in the clear or to any other destination than the ASA. This, in effect, disables split tunneling. Remote users reach Internet networks through the corporate network and do not have access to local networks. This is the default option.

Step 7 In the **Network List** field, select the unified access control rule to which you apply the split-tunneling policy. If you do not select Inherit, the group policy uses the network list specified in the default group policy. The default value of the network list in the default group policy is --None--.

Use the **Manage** command button to open the ACL Manager dialog box, in which you can configure access control lists, or select an existing access control list to use as network lists. If you want to create or edit a network list, see Adding ACLs and ACEs, page 26-2.

Note Extended ACL lists can be used for both IPv4 and IPv6 addresses.

Step 8 The **Intercept DHCP Configuration Message from Microsoft Clients** reveals additional parameters specific to DHCP Intercept. DHCP Intercept lets Microsoft XP clients use split-tunneling with the ASA. 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.

Step 9 Click OK.After making this change to the group policy, be sure the group policy is associated with a Connection Profile used by AnyConnect. In ASDM select, **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles** to configure a Connection Profile.

Configuring a Browser Proxy for an Internal Group Policy

Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > Advanced > 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.

- Proxy Auto Configuration Settings—The 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.

Configuring General AnyConnect Client Attributes for an Internal Group Policy

Clicking the AnyConnect Client icon in the group policy directory tree shows the list of configurable attributes that follow. Configuring the ASA to distribute and manage AnyConnect client sessions is a larger procedure than just setting these attribute fields in a group policy. See Configuring AnyConnect VPN Client Connections, page 87-42, Configuring AnyConnect VPN Connections, page 87-52, and Configuring AnyConnect Secure Mobility, page 87-64.

Fields

- Keep Installer on Client System—Enable 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.

Note Keep Installer on Client System is not supported after version 2.5 of the AnyConnect client.

- Datagram Transport Layer Security (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 Don't Defrag (DF) Bit—This feature allows the force fragmentation of packets that have the DF bit set, allowing them to pass through the tunnel. An example use case is for servers in your network that do not respond correctly to TCP MSS negotiations.
- Client Bypass Protocol—The Client Protocol Bypass feature allows you to configure how the ASA manages IPv4 traffic when it is expecting only IPv6 traffic or how it manages IPv6 traffic when it is expecting only IPv4 traffic.

  When the AnyConnect client makes a VPN connection to the ASA, the ASA could assign it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the ASA assigns the AnyConnect connection only an IPv4 address or only an IPv6 address, you can now configure the Client Bypass Protocol to drop network traffic for which the ASA did not assign an IP address, or allow that traffic to bypass the ASA and be sent from the client unencrypted or "in the clear".

  For example, assume that the ASA assigns only an IPv4 address to an AnyConnect connection and the endpoint is dual stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.
- FQDN of This Device—This information is used by the client after network roaming in order to resolve the ASA IP address used for re-establishing the VPN session. This setting is critical to support roaming between networks of different IP protocols (such as IPv4 to IPv6).

  Note You cannot use the ASA FQDN present in the AnyConnect profile to derive the ASA IP address after roaming. The addresses may not match the correct device (the one the tunnel was established to) in the load balancing scenario.

  If the device FQDN is not pushed to the client, the client will try to reconnect to whatever IP address the tunnel had previously established. In order to support roaming between networks of different IP protocols (from IPv4 to IPv6), AnyConnect must perform name resolution of the device FQDN after roaming, so that it can determine which ASA address to use for re-establishing the tunnel. The client uses the ASA FQDN present in its profile during the initial connection. During subsequent session reconnects, it always uses the device FQDN pushed by ASA (and configured by the administrator in the group policy), when available. If the FQDN is not configured, the ASA derives the device FQDN (and sends it to the client) from whatever is set under Device Setup > Device Name/Password and Domain Name.

  If the device FQDN is not pushed by the ASA, the client cannot re-establish the VPN session after roaming between networks of different IP protocols.
- 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.
- Keepalive Messages—Enter a 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.
- **Optional Client Modules to Download**—To minimize download time, the AnyConnect client requests downloads (from the ASA) only of modules that it needs for each feature that it supports. You must specify the names of modules that enable other features. The AnyConnect client, version 3.0, includes the following modules (previous versions have fewer modules):

- AnyConnect DART—The Diagnostic AnyConnect Reporting Tool (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.
- AnyConnect Network Access Manager—Formerly called the Cisco Secure Services Client, this module provides 802.1X (Layer 2) and device authentication for access to both wired and wireless network is integrated into AnyConnect 3.0.
- AnyConnect SBL—Start Before Logon (SBL) forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears.
- AnyConnect Web Security Module—Formerly called ScanSafe Hostscan, this module is integrated into the AnyConnect 3.0.
- AnyConnect Telemetry Module—Sends information about the origin of malicious content to the web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA), which uses this data to provide better URL filtering rules.
- AnyConnect Posture Module—Formerly called the Cisco Secure Desktop HostScan feature, the posture module is integrated into AnyConnect 3.0 and provides AnyConnect the ability to gather credentials for posture assessment prior to creating a remote access connection to the ASA.

- Always-On VPN—Determine if the always-on VPN flag setting in the AnyConnect service profile is disabled or if the AnyConnect service profile setting should be used. The always-on VPN feature lets AnyConnnect automatically establish a VPN session after the user logs onto a computer. The VPN session remains up until the user logs off the computer. If the physical connection is lost, the session remains up, and AnyConnect continually attempts to reestablish the physical connection with the adaptive security appliance to resume the VPN session.

  Always-on VPN permits the enforcement of corporate policies to protect the device from security threats. You can use it to help ensure AnyConnect establishes a VPN session whenever the endpoint is not in a trusted network. If enabled, a policy is configured to determine how network connectivity is managed in the absence of a connection.

  Note Always-On VPN requires an AnyConnect release that supports AnyConnect Secure Mobility features. Refer to the *Cisco AnyConnect VPN Client Administrator Guide* for additional information.

- Client Profiles to Download—A profile is a group of configuration parameters that the AnyConnect client uses to configure VPN, Network Access Manager, web security, and telemetry settings. Click Add to launch the Select Anyconnect Client Profiles window where you can specify previously-created profiles for this group policy.

Configuring AnyConnect Login Settings for an Internal Group Policy

In this dialog box, you can enable the ASA to prompt remote users to download the AnyConnect client or go to a Clientless SSL VPN portal page. Figure 87-1 shows the prompt displayed:

Figure 87-1 Prompt Displayed to Remote Users for AnyConnect Client Download

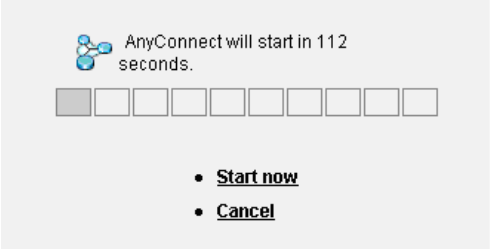

Fields

- 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.

Configuring AnyConnect Client Client Firewall Attributes for an Internal Group Policy

In ASA 9.0 and later releases, access control lists for client firewalls support both access list entries for both IPv4 and IPv6 addresses.

See Client Firewall with Local Printer and Tethered Device Support, page 87-37 to configure the group policy for these situations.

Configuring AnyConnect Client Key Regeneration for an Internal Group Policy

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—Uncheck the Inherit check box to specify a renegotiation method different from the default group policy. Select the **None** radio button to disable rekey, select either the **SSL** or **New Tunnel** radio button to establish a new tunnel during rekey.

Note Configuring the Renegotiation Method as **SSL** or **New Tunnel** specifies that the client establishes a new tunnel during rekey instead of the SSL renegotiation taking place during the rekey. See the Cisco ASA 5500 Series Command Reference, 8.4 for a history of the **anyconnect ssl rekey** command.

Configuring AnyConnect Client Dead Peer Detection for an Internal Group Policy

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.

Customizing a VPN Access Portal for an Internal 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—Selects the customization to apply to the AnyConnect Client/SSL VPN portal page. The default is DfltCustomization.
  - 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—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.
- 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**.

  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.”

Configuring AnyConnect Client Custom Attributes for an Internal Group Policy

This dialog box lists the custom attributes that are assigned to this group policy. Custom attributes can be created in this dialog, or on Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Custom Attributes. In this dialog, you can add custom attributes to this group policy, and define values for those attributes.

For AnyConnect 3.1, custom attributes are available to support AnyConnect Deferred Upgrade and Phone Home.

See the *Cisco AnyConnect Secure Mobility Client Administrator Guide, Release 3.1*, Chapter 2, “User Control Over Upgrade” for more information about the deferred upgrade custom attributes.

IPsec (IKEv1) Client

Configuring IPsec (IKEv1) Client General Attributes for an Internal Group Policy

Configuration > Remote Access > Network (Client) Access > Group Policies > Advanced > IPsec (IKEv1) 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—Allows 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.

Configuring IPsec (IKEv1) Client Client Access Rules for an Internal Group Policy

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.
- 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.
- 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.

Configuring IPsec (IKEv1) Client Client Firewall Attributes for an Internal Group Policy

Configuration > Remote Access > Network (Client) Access > Group Policies > Advanced > IPsec (IKEv1) Client > Client Firewall Tab

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 ASA 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 ASA. (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 ASA.) 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 ASA, 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 ASA 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 ASA drops any session that attempts to connect without the designated, supported firewall installed and running. In this case, the ASA 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 ASA 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 ASA 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 ASA, including the default filters. Keep in mind that the ASA pushes these rules down to the VPN client, so you should create and define these rules relative to the VPN client, not the ASA. 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 ASA 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).

Configuring IPsec (IKEv1) Client Hardware Client Attributes for an Internal Group Policy

Configuration > Remote Access > Network (Client) Access > Group Policies > Advanced > IPsec (IKEv1) Client > 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 ASA to which it connects. The ASA 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 ASA 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 ASA, 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 ASA 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 ASA 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 ASA, 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 ASA, or if you direct the browser to a website on the remote network behind the ASA, 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 ASA 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 ASA 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 ASA 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 ASA 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:

Configuring Clientless SSL VPN Internal Group Policies

Configuring Clientless SSL VPN General Attributes for an Internal Group Policy

Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies > Add/Edit > Add or Edit Internal Group Policy > General

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.

  The IPsec VPN client supports full HTML for the banner. However, the clientless portal and the AnyConnect client support partial HTML. To ensure the banner displays properly to remote users, follow these guidelines:

  - For IPsec client users, use the /n tag.
  - For AnyConnect client users, use the <BR> tag.
  - For clientless users, use the <BR> tag.
- 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 an ASA; 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. If you are using the AnyConnect client, you must choose this protocol for MUS to be supported.
  - IPsec IKEv1—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 Cisco VPN client-to-LAN connections can use IPsec IKEv1.
  - IPsec IKEv2—Supported by the AnyConnect Secure Mobility Client. AnyConnect connections using IPsec with IKEv2 provide advanced features such as software updates, client profiles, GUI localization (translation) and customization, Cisco Secure Desktop, and SCEP proxy.
  - 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.
- 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.

- 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. See Defining Time Ranges, page 87-29 for more information.
- 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 ASA 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 ASA.

Note This feature works for HTTP connections, but not for FTP and CIFS.

- Connection Profile (Tunnel Group) Lock—This parameter permits remote VPN access only with the selected connection profile (tunnel group), and prevents access with a different connection profile. The default inherited value is None.
- 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.
- Session Alert Interval— If you uncheck the Inherit check box, the Default checkbox is checked automatically. This sets the session alert interval to 30 minutes. If you want to specify a new value, uncheck the Default check box and specify a session alert interval from 1 to 30 minutes in the minutes box.
- Idle Alert Interval—If you uncheck the Inherit check box, the Default checkbox is checked automatically. This sets the idle alert interval to 30 minutes. If you want to specify a new value, uncheck the Default check box and specify a session alert interval from 1 to 30 minutes in the minutes box.

Configuring the Clientless SSL VPN Access Portal for an Internal 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—Specify your smart tunnel options using a clientless (browser-based) SSL VPN session with the ASA as the pathway and the security appliance as a proxy server:
  - Smart Tunnel Policy—Choose from the network list and specify one of the tunnels options: use smart tunnel for the specified network, do not use smart tunnel for the specified network, or use tunnel for all network traffic. Assigning a smart tunnel network 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**.
  - Smart Tunnel Application—Choose from the drop-down menu to connect a Winsock 2, TCP-based application installed on the end station to a server on the intranet. To view, add, modify, or delete a smart tunnel application, click **Manage**.
  - Smart Tunnel all Applications—Check this check box to tunnel all applications. All applications are tunneled without choosing from the network list or knowing which executables an end user may invoke for external applications.
  - Auto Start—Check this check box to start smart tunnel access automatically upon user login. This option to start smart tunnel access upon user login applies only to Windows. 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.

Configuring Portal Customization for a Clientless SSL VPN Internal 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.

The procedure for customizing an access portal for a Clientless SSL VPN Access connection is the same as it is for a Network Client Access connection. See Customizing a VPN Access Portal for an Internal Group Policy, page 87-17.

Configuring Login Settings for a Clientless SSL VPN Internal Group Policy

In this dialog box, you can enable the ASA to prompt remote users to download the AnyConnect client or go to a Clientless SSL VPN portal page. See Configuring AnyConnect Login Settings for an Internal Group Policy, page 87-15.

Configuring Single Signon and Auto Signon Servers for a Clientless SSL VPN Access Internal Group Policy

To configure single sign-on servers and Auto sign-on servers, see "Auto Signon" in Chapter 71, "Clientless SSL VPN" in *Cisco ASA 5500 Series Configuration Guide using ASDM, 6.4 and 6.6* or "Using Single Sign-on with Clientless SSL VPN" in Chapter 73 of *Cisco ASA 5500 Series Configuration Guide using the CLI, 8.4 and 8.6.*

Configuring Session Settings for Clientless SSL VPN Access

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. See “Configuring Session Settings” Chapter 71, “Clientless SSL VPN” in *Cisco ASA 5500 Series Configuration Guide using ASDM, 6.4 and 6.6* or in Chapter 73 of *Cisco ASA 5500 Series Configuration Guide using the CLI, 8.4 and 8.6.*

Configuring Site-to-Site Internal Group Policies

Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > Add or Edit Internal Group Policy > General

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 ASA; 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. If you are using the AnyConnect client, you must choose this protocol for MUS to be supported.
  - IPsec IKEv1—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 Cisco VPN client-to-LAN connections can use IPsec IKEv1.
  - IPsec IKEv2—Supported by the AnyConnect Secure Mobility Client. AnyConnect connections using IPsec with IKEv2 provide advanced features such as software updates, client profiles, GUI localization (translation) and customization, Cisco Secure Desktop, and SCEP proxy.
  - 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.

- 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 ASA, 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.
- 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.
- 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).

Defining Time Ranges

You can get to this panel through various paths.

Use the Browse Time Range dialog box to add, edit, or delete a time range. A time range is a reusable component that defines starting and ending times that can be applied to a group policy. After defining a time range, you can select the time range and apply it to different options that require scheduling. For example, you can attach an access list to a time range to restrict access to the ASA. 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.

Add/Edit Time Range

You can get to this panel through various paths.

The Add or Edit Time Range dialog box lets you configure a new time range.

Fields

- Time Range Name—Specifies the name that you want to assign to this time range.

- Start Time—Defines the time when you want the time range to start.
  - Start now—Specifies that the time range starts immediately.
  - Start at—Selects the month, day, year, hour, and minute at which you want the time range to start.
- End Time—Defines the time when you want the time range to end.
  - Never end—Specifies that the time range has no defined end point.
  - End at (inclusive)—Selects the month, day, year, hour, and minute at which you want the time range to end.
- Recurring Time Ranges—Constrains the active time of this time range within the start and end times when the time range is active. For example, if the start time is start now and the end time is never end, and you want the time range to be effective every weekday, Monday through Friday, from 8:00 AM to 5:00 PM, you could configure a recurring time range, specifying that it is to be active weekdays from 08:00 through 17:00, inclusive.
- Add—Opens the Add Recurring Time Range dialog box, 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.

Add/Edit Recurring Time Range

You can get to this panel through various paths.

The Add or Edit Recurring Time Range dialog box lets you configure or modify a recurring time range.

Fields

- Specify days of the week and times on which this recurring range will be active—Makes available the options in the Days of the week area. For example, use this option when you want the time range to be active only every Monday through Thursday, from 08:00 through 16:59.
  - Days of the week—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.

Access Control List Manager

You can get to this panel through various paths.

The ACL Manager dialog box lets you define access control lists (ACLs) to control the access of a specific host or network to another host/network, including the protocol or port that can be used.

You can configure ACLs (access control lists) to apply to user sessions. These are filters that permit or deny user access to specific networks, subnets, hosts, and web servers.

- If you do not define any filters, all connections are permitted.
- The ASA 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 ASA denies it. ACEs are referred to as rules in this section.

Standard Access Control List

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.”

Extended Access Control List

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.”

Add/Edit/Paste ACE

ACL Manager > Add/Edit/Paste Extended ACE

The Add/Edit/Paste ACE dialog box lets you create a new extended access list rule, or modify an existing rule. The Paste option becomes available only when you cut or copy a rule.

Fields

- Action—Determines the action type of the new rule. Select either permit or deny.
  - Permit—Permits all matching traffic.
  - Deny—Denies all matching traffic.

- Source/Destination—Specifies the source or destination type and, depending on that type, the other relevant parameters describing the source or destination host/network IP Address. Possible values are: any, IP address, Network Object Group, and Interface IP. The availability of subsequent fields depends upon the value of the Type field:
  - any—Specifies that the source or destination host/network can be any type. For this value of the Type field, there are no additional fields in the Source or Destination area.
  - IP Address—Specifies the source or destination host or network IP address. 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.

Browse Source/Destination Address

ACL Manager > Add/Edit Extended Access List Rule > Source or Destination > Browse button

The Browse Source or Destination Address dialog box lets you select an object to use 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.

Browse Source/Destination Port

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: tcp or udp >Source or Destination Port > Group option > Browse button

The Browse Source or Destination Port dialog box lets you select a source or destination port for this protocol in this rule.

Fields

- Add—Opens the Add TCP Service Group dialog box, 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.

Add TCP Service Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: tcp or udp >Source or Destination Port > Group option > Browse button > Browse Source or Destination Port > Add button

The Add TCP Service Group dialog box lets you configure a new a TCP service group or port to add to the browsable source or destination port list for this protocol in this rule. Selecting a member of either the Members not in Group or the Members in Group list activates the Add and Remove buttons.

Fields

- Group Name—Specifies the name of the new TCP service group.
- Description—(Optional) Provides a brief description of this group.
- Members not in Group—Presents the option to select either a service/service group or a port number to add to the Members in Group list.
- Service/Service Group—Selects the option to select the name of a TCP service or service group to add to the Members in Group list.
- Port #—Selects the option to specify a range of port numbers to add to the Members in Group list.
- Add—Moves a selected item from the Members not in Group list to the Members in Group list.
- Remove—Moves a selected item from the Members in Group list to the Members not in Group list.
- Members in Group—Lists the members already configured in this service group.

Browse ICMP

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: icmp >ICMP > Group option > Browse button

The Browse ICMP dialog box lets you select an ICMP group for this rule.

Fields

- Add—Opens the Add ICMP Group dialog box, 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.

Add ICMP Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: icmp >ICMP > Group option > Browse button > Browse ICMP > Add button

The Add ICMP Group dialog box lets you configure a new a ICMP group by name or by number to add to the browsable ICMP list for this protocol in this rule. 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.

Browse Other

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: other >Other > Group option > Browse button

The Browse Other dialog box lets you select a protocol group for this rule.

Fields

- Add—Opens the Add Protocol Group dialog box, 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.

Add Protocol Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: other > Group option > Browse button > Browse Other > Add button

The Add Protocol Group dialog box lets you configure a new a protocol group by name or by number to add to the browsable protocol list for this rule. Selecting a member of either the Members not in Group or the Members in Group list activates the Add and Remove buttons.

Fields

- Group Name—Specifies the name of the new TCP service group.
- Description—(Optional) Provides a brief description of this group.
- Members not in Group—Presents the option to select either a protocol/protocol group or a protocol number to add to the Members in Group list.
- Protocol/Protocol Group—Selects the option to select the name of a protocol or protocol group to add to the Members in Group list.
- Protocol #—Selects the option to specify a protocol by number to add to the Members in Group list.
- Add—Moves a selected item from the Members not in Group list to the Members in Group list.
- Remove—Moves a selected item from the Members in Group list to the Members not in Group list.
- Members in Group—Lists the members already configured in this service group.

Client Firewall with Local Printer and Tethered Device Support

When users connect to the ASA, all traffic is tunneled through the connection and users cannot access resources on their local network. This includes printers, cameras, and Windows Mobile devices (tethered devices) that synchronize with the local computer. Enabling Local LAN Access in the client profile resolves this problem, however it can introduce a security or policy concern for some enterprises as a result of unrestricted access to the local network. You can use the ASA to deploy endpoint OS firewall capabilities to restrict access to particular types of local resources, such as printers and tethered devices.

To do so, enable client firewall rules for specific ports for printing. The client distinguishes between inbound and outbound rules. For printing capabilities, the client opens ports required for outbound connections, but blocks all incoming traffic.

Note Be aware that users logged in as administrators have the ability to modify the firewall rules deployed to the client by the ASA. Users with limited privileges cannot modify the rules. For either user, the client reapplies the firewall rules when the connection terminates.

If you configure the client firewall, and the user authenticates to an Active Directory (AD) server, the client still applies the firewall policies from the ASA. However, the rules defined in the AD group policy take precedence over the rules of the client firewall.

The following sections describe procedures on how to do this:

- Deploying a Client Firewall for Local Printer Support, page 87-38
- Tethered Devices Support, page 87-39

Usage Notes about Firewall Behavior

The following notes clarify how the AnyConnect client uses the firewall:

- The source IP is not used for firewall rules. The client ignores the source IP information in the firewall rules sent from the ASA. The client determines the source IP depending on whether the rules are public or private. Public rules are applied to all interfaces on the client. Private rules are applied to the Virtual Adapter.

- The ASA supports many protocols for ACL rules. However, the AnyConnect firewall feature supports only TCP, UDP, ICMP, and IP. If the client receives a rule with a different protocol, it treats it as an invalid firewall rule, and then disables split tunneling and uses full tunneling for security reasons.
- Starting in ASA 9.0, the Public Network Rule and Private Network Rule support unified access control lists. These access control lists can be used to define IPv4 and IPv6 traffic in the same rule.

Be aware of the following differences in behavior for each operating system:

- For Windows computers, deny rules take precedence over allow rules in Windows Firewall. If the ASA pushes down an allow rule to the AnyConnect client, but the user has created a custom deny rule, the AnyConnect rule is not enforced.
- On Windows Vista, when a firewall rule is created, Vista takes the port number range as a comma-separated string. The port range can be a maximum of 300 ports. For example, from 1-300 or 5000-5300. If you specify a range greater than 300 ports, the firewall rule is applied only to the first 300 ports.
- Windows users whose firewall service must be started by the AnyConnect client (not started automatically by the system) may experience a noticeable increase in the time it takes to establish a VPN connection.
- On Mac computers, the AnyConnect client applies rules sequentially in the same order the ASA applies them. Global rules should always be last.
- For third-party firewalls, traffic is passed only if both the AnyConnect client firewall and the third-party firewall allow that traffic type. If the third-party firewall blocks a specific traffic type that the AnyConnect client allows, the client blocks the traffic.

Deploying a Client Firewall for Local Printer Support

The ASA supports the AnyConnect client firewall feature with ASA version 8.3(1) or later, and ASDM version 6.3(1) or later. This section describes how to configure the client firewall to allow access to local printers, and how to configure the client profile to use the firewall when the VPN connection fails.

Limitations and Restrictions of the Client Firewall

The following limitations and restrictions apply to using the client firewall to restrict local LAN access:

- Due to limitations of the OS, the client firewall policy on computers running Windows XP is enforced for inbound traffic only. Outbound rules and bidirectional rules are ignored. This would include firewall rules such as 'permit ip any any'.
- Host Scan and some third-party firewalls can interfere with the firewall.

Table 72-1 clarifies what direction of traffic is affected by the source and destination port settings:

Source Port	Destination Port	Traffic Direction Affected
Specific port number	Specific port number	Inbound and outbound
A range or 'All' (value of 0)	A range or 'All' (value of 0)	Inbound and outbound
Specific port number	A range or 'All' (value of 0)	Inbound only
A range or 'All' (value of 0)	Specific port number	Outbound only

Example ACL Rules for Local Printing

Table 87-1 shows example ACL rules for local printing:

Table 87-1 Example ACL Rules for Local Printing

Description	Permission	Interface	Protocol	Source Port	Destination Address	Destination Port
Deny all	Deny	Public	Any	Default[1]	Any	Default
LPD	Allow	Public	TCP	Default	Any	515
IPP	Allow	Public	TCP	Default	Any	631
Printer	Allow	Public	TCP	Default	Any	9100
mDNS	Allow	Public	UDP	Default	224.0.0.251	5353
LLMNR	Allow	Public	UDP	Default	224.0.0.252	5355
NetBios	Allow	Public	TCP	Default	Any	137
NetBios	Allow	Public	UDP	Default	Any	137

1. The port range is 1 to 65535.

Note To enable local printing, you must enable the **Local LAN Access** feature in the client profile with a defined ACL rule *allow Any Any.*

Configuring Local Print Support

Step 1 Enable the AnyConnect client firewall in a group policy. Go to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**.

Step 2 Select a group policy and click **Edit**. The Edit Internal Group Policy window displays.

Step 3 Select **Advanced > AnyConnect Client > Client Firewall**. Click **Manage** for the Private Network Rule.

Step 4 Create an ACL and specify an ACE using the rules in Table 87-1. Add this ACL as a PrivateNetwork Rule.

Step 5 If you enabled the Automatic VPN Policy always-on and specified a closed policy, in the event of a VPN failure, users have no access to local resources. You can apply the firewall rules in this scenario by going to **Preferences (Cont) in the profile editor and** checking **Apply last local VPN resource rules**.

Tethered Devices Support

To support tethered devices and protect the corporate network, create a standard ACL in the group policy, specifying destination addresses in the range that the tethered devices use. Then specify the ACL for split tunneling as a network list to exclude from tunneled VPN traffic. You must also configure the client profile to use the last VPN local resource rules in case of VPN failure.

Note For Windows Mobile devices that need to sync with the computer running AnyConnect, specify the IPv4 destination address as 169.254.0.0, or the IPv6 destination address fe80::/64 in the ACL.

Follow these steps:

Step 1 In ASDM, go to **Group Policy > Advanced > Split Tunneling**.

Step 2 Uncheck **Inherit** next to the Network List field and click **Manage**. The ACL Manager displays.

Step 3 Click the **Extended ACL** tab.

Step 4 Click **Add** and then **Add ACL**. Specify a name for the new ACL.

Step 5 Choose the new ACL in the table and click **Add** and then **Add ACE**. The Edit ACE window displays.

Step 6 For Action, choose the **Permit** radio button.

Step 7 In the destination criteria area, specify the IPv4 destination address as 169.254.0.0 or the IPv6 destination address fe80::/64.

Step 8 For Service, choose *IP*.

Step 9 Click **OK**.

Step 10 Click OK to save the ACL.

Step 11 In the Split Tunneling pane for the internal group policy, uncheck Inherit for the Policy or IPv6 Policy, depending on the IP address you specified in step 7, and choose **Exclude Network List Below**. For Network List, choose the ACL you created.

Step 12 Click **OK**.

Step 13 Click **Apply**.

Fields

- Public Network Rule—Use the drop-down list to choose from the existing defined ACLs.

  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).

- Private Network Rule—Use the drop-down list to choose from the existing defined ACLs.

  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).

Configuring a Web ACLs

Configuration > Remote Access > Clientless SSL VPN Access > Advanced > Web ACLs

Configuration > Remote Access > Clientless SSL VPN Access > Group Policies > General > More Options > Web ACL

This dialog box lets you configure ACLs for Clientless SSL VPN connections.

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 ASA checks ACLs and their ACEs in priority order according to their position in the ACLs list box until it finds a match.

Add/Edit Standard Access List Rule

ACL Manager > Add or Edit Standard Access List Rule

The Add/Edit Standard Access List Rule dialog box lets you create a new rule, or modify an existing rule.

Fields

- Action—Determines the action type of the new rule. 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.

Add/Edit Server and URL List

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Other Tab > Add or Edit Server and URL List

The Add or Edit Server and URL List dialog box lets you add, edit, delete, and order the items in the designated URL list.

Fields

- List Name—Specifies the name of the list to be added or selects the name of the list to be modified or deleted.
- URL Display Name—Specifies the URL name displayed to the user.
- URL—Specifies the actual URL associated with the display name.

- Add—Opens the Add Server or URL dialog box, 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

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Other Tab > Add or Edit Server and URL

The Add or Edit Server or URL dialog box lets you add or edit, delete, and order the items in the designated URL list.

Fields

- URL Display Name—Specifies the URL name displayed to the user.
- URL—Specifies the actual URL associated with the display name.

Configuring AnyConnect VPN Client Connections

The Cisco AnyConnect VPN client provides secure SSL or IPsec (IKEv2) connections to the ASA for remote users. The client gives remote users the benefits of a 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 ASA 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 ASA 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 VPN connection and either remains or uninstalls itself (depending on the ASA configuration) when the connection terminates.

In the case of a previously installed client, when the user authenticates, the ASA examines the revision of the client, and upgrades the client as necessary.

The AnyConnect client can be downloaded from the ASA, or it can be installed manually on the remote PC by the system administrator. For more information about installing the client manually, see the *AnyConnect Administrators Guide*.

The ASA downloads the client based on the group policy or username attributes of the user establishing the connection. You can configure the ASA 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 ASA 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.

Note Keep Installer on Client System is not supported after version 2.5 of the AnyConnect client.

- Compression—Compression increases the communications performance between the security appliance and the client by reducing the size of the packets being transferred.
- Datagram TLS—Datagram Transport Layer Security 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 Don't Defrag (DF) Bit—This feature allows the force fragmentation of packets that have the DF bit set, allowing them to pass through the tunnel. An example use case is for servers in your network that do not respond correctly to TCP MSS negotiations.
- Client Bypass Protocol—The Client Protocol Bypass feature allows you to configure how the ASA manages IPv4 traffic when it is expecting only IPv6 traffic or how it manages IPv6 traffic when it is expecting only IPv4 traffic.

  When the AnyConnect client makes a VPN connection to the ASA, the ASA could assign it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the ASA assigns the AnyConnect connection only an IPv4 address or only an IPv6 address, you can now configure the Client Bypass Protocol to drop network traffic for which the ASA did not assign an IP address, or allow that traffic to bypass the ASA and be sent from the client unencrypted or "in the clear".

  For example, assume that the ASA assigns only an IPv4 address to an AnyConnect connection and the endpoint is dual stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.
- FQDN of This Device—This information is used by the client after network roaming in order to resolve the ASA IP address used for re-establishing the VPN session. This setting is critical to support roaming between networks of different IP protocols (such as IPv4 to IPv6).

  Note You cannot use the ASA FQDN present in the AnyConnect profile to derive the ASA IP address after roaming. The addresses may not match the correct device (the one the tunnel was established to) in the load balancing scenario.

  If the device FQDN is not pushed to the client, the client will try to reconnect to whatever IP address the tunnel had previously established. In order to support roaming between networks of different IP protocols (from IPv4 to IPv6), AnyConnect must perform name resolution of the device FQDN after roaming, so that it can determine which ASA address to use for re-establishing the tunnel. The client uses the ASA FQDN present in its profile during the initial connection. During subsequent session reconnects, it always uses the device FQDN pushed by ASA (and configured by the administrator in the group policy), when available. If the FQDN is not configured, the ASA derives the device FQDN (and sends it to the client) from whatever is set under Device Setup > Device Name/Password and Domain Name.

  If the device FQDN is not pushed by the ASA, the client cannot re-establish the VPN session after roaming between networks of different IP protocols.

- 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.
- Keepalive Messages—Enter a 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.
- **Optional Client Modules to Download**—To minimize download time, the AnyConnect client requests downloads (from the ASA) only of modules that it needs for each feature that it supports. You must specify the names of modules that enable other features. The AnyConnect client, version 3.0, includes the following modules (previous versions have fewer modules):
  - AnyConnect DART—The Diagnostic AnyConnect Reporting Tool (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.
  - AnyConnect Network Access Manager—Formerly called the Cisco Secure Services Client, this module provides 802.1X (Layer 2) and device authentication for access to both wired and wireless network is integrated into AnyConnect 3.0.
  - AnyConnect SBL—Start Before Logon (SBL) forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears.
  - AnyConnect Web Security Module—Formerly called ScanSafe Hostscan, this module is integrated into the AnyConnect 3.0.
  - AnyConnect Telemetry Module—Sends information about the origin of malicious content to the web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA), which uses this data to provide better URL filtering rules.
  - AnyConnect Posture Module—Formerly called the Cisco Secure Desktop HostScan feature, the posture module is integrated into AnyConnect 3.0 and provides AnyConnect the ability to gather credentials for posture assessment prior to creating a remote access connection to the ASA.
- Always-On VPN—Determine if the always-on VPN flag setting in the AnyConnect service profile is disabled or if the AnyConnect service profile setting should be used. The always-on VPN feature lets AnyConnnect automatically establish a VPN session after the user logs onto a computer. The VPN session remains up until the user logs off the computer. If the physical connection is lost, the session remains up, and AnyConnect continually attempts to reestablish the physical connection with the adaptive security appliance to resume the VPN session.

  Always-on VPN permits the enforcement of corporate policies to protect the device from security threats. You can use it to help ensure AnyConnect establishes a VPN session whenever the endpoint is not in a trusted network. If enabled, a policy is configured to determine how network connectivity is managed in the absence of a connection.

Note Always-On VPN requires an AnyConnect release that supports AnyConnect Secure Mobility features. Refer to the *Cisco AnyConnect VPN Client Administrator Guide* for additional information.

- Client Profiles to Download—A profile is a group of configuration parameters that the AnyConnect client uses to configure VPN, Network Access Manager, web security, and telemetry settings. Click Add to launch the Select Anyconnect Client Profiles window where you can specify previously-created profiles for this group policy.

Using AnyConnect Client Profiles

You enable Cisco AnyConnect Secure Mobility client features in the AnyConnect profiles—XML files that contain configuration settings for the core client with its VPN functionality and for the optional client modules Network Access Manager, telemetry, and web security. The ASA deploys the profiles during AnyConnect installation and updates. Users cannot manage or modify profiles.

You can configure a profile using the AnyConnect profile editor, a convenient GUI-based configuration tool launched from ASDM. The AnyConnect software package, version 2.5 and later (for all OSs), includes the editor, which activates when you load the AnyConnect package on the ASA as an AnyConnect client image. Alternatively, you can manually edit the XML file and import the file to the ASA as a profile.

You can configure the ASA to deploy profiles globally for all AnyConnect users or to users based on their group policy. Usually, a user has a single profile file for each AnyConnect module installed. In some cases, you might want to provide more than one profile for a user. Someone who works from multiple locations might need more than one profile. Be aware that some of the profile settings (such as SBL) control the connection experience at a global level. Other settings are unique to a particular host and depend on the host selected.

Some profile settings are stored locally on the user computer in a user preferences file or a global preferences file. The user file has information the client needs to display user-controllable settings in the Preferences tab of the client GUI and information about the last connection, such as the user, the group, and the host. The global file has information about user-controllable settings to be able to apply those settings before login (since there is no user). For example, the client needs to know if Start Before Logon and/or AutoConnect On Start are enabled before login. For more information about creating and deploying AnyConnect client profiles and controlling client features, see the *AnyConnect VPN Client Administrator Guide*.

Fields

Add—Displays the Add AnyConnect 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 Profile window, where you can change the settings contained in the profile for AnyConnect client features.

Delete—Deletes a profile from the table. This does not delete the XML file from flash.

AnyConnect Client Profiles Table—Displays the XML files specified as AnyConnect client profiles:

- **Profile Name**—The name of the profile specified when the profile was added.
- **Profile Usage/Type**—Displays the use for this profile, such as VPN, Network Access Manager, or telemetry.

Specifying an AnyConnect Client Profile

Specify an AnyConnect client profile for this group policy.

For more information about creating and deploying AnyConnect client profiles and controlling client features, see the *AnyConnect VPN Client Administrator Guide*.

Fields

Profile Name—Specify an AnyConnect client profile for this group policy.

Profile Usage—Displays the usage assigned to the profile when originally created: VPN, Network Access Manager, web security, or telemetry. If ASDM does not recognize the usage specified in the XML file, the drop-down list becomes selectable and you can choose a usage type manually.

Profile Location—Specify a path to the profile file in the ASA flash memory. If the file does not exist, the ASA creates one based on the profile template.

Importing an AnyConnect Client Profile

Import a new AnyConnect client profile in this window. You can import a profile from a local device or a remote server.

For more information about creating and deploying AnyConnect client profiles and controlling client features, see the *AnyConnect VPN Client Administrator Guide*.

Fields

Profile Name—Specify a name for the profile you add.

Profile Usage—Displays the usage assigned to the profile when originally created: VPN, Network Access Manager, web security, or telemetry. If ASDM does not recognize the usage specified in the XML file, the drop-down list becomes selectable and you can choose a usage type manually.

Group Policy—Specify a group policy for this profile. The profile downloads to users belonging to the group policy along with the AnyConnect client.

Profile Location—Specify a path to the profile file in the ASA flash memory. If the file does not exist, the ASA creates one based on the profile template.

Exporting an AnyConnect Client Profile

Export an AnyConnect VPN client profile from this window. You can export to a local device or a remote server.

For more information about creating and deploying AnyConnect client profiles and controlling client features, see the *AnyConnect VPN Client Administrator Guide*.

Fields

Device Profile Path—Displays the path and filename of the profile file.

Local Path—Specify the path and filename to export the profile file.

Browse Local—Click to launch a window to browse the local device file system.

Exempting AnyConnect Traffic from Network Address Translation

If you have configured your ASA to perform network address translation (NAT), you must exempt your remote access AnyConnect client traffic from being translated so that the AnyConnect clients, internal networks, and corporate resources on a DMZ, can originate network connections to each other. Failing to exempt the AnyConnect client traffic from being translated prevents the AnyConnect clients and other corporate resources from communicating.

"Identity NAT" (also known as "NAT exemption") allows an address to be translated to itself, which effectively bypasses NAT. Identity NAT can be applied between two address pools, an address pool and a subnetwork, or two subnetworks.

This procedure illustrates how you would configure identity NAT between these hypothetical network objects in our example network topology: Engineering VPN address pool, Sales VPN address pool, inside network, a DMZ network, and the Internet. Each Identity NAT configuration requires one NAT rule.

Table 87-2 Network Addressing for Configuring Identity NAT for VPN Clients

Network or Address Pool	Network or address pool name	Range of addresses
Inside network	inside-network	10.50.50.0 - 10.50.50.255
Engineering VPN address pool	Engineering-VPN	10.60.60.1 - 10.60.60.254
Sales VPN address pool	Sales-VPN	10.70.70.1 - 10.70.70.254
DMZ network	DMZ-network	192.168.1.0 - 192.168.1.255

Step 1 Log into the ASDM and select **Configuration > Firewall > NAT Rules**.

Step 2 Create a NAT rule so that the hosts in the Engineering VPN address pool can reach the hosts in the Sales VPN address pool. In the NAT Rules pane, select **Add > Add NAT Rule Before "Network Object" NAT rules** so that the ASA evaluates this rule before other rules in the Unified NAT table. See Figure 87-2 on page 87-48 for an example of the Add NAT rule dialog box.

Note NAT rule evaluation is applied on a top-down, first match basis. Once the ASA matches a packet to a particular NAT rule it does not perform any further evaluation. It is important that you place the most specific NAT rules at the top of the Unified NAT table so that the ASA does not prematurely match them to broader NAT rules.

Figure 87-2 Add NAT rule dialog box

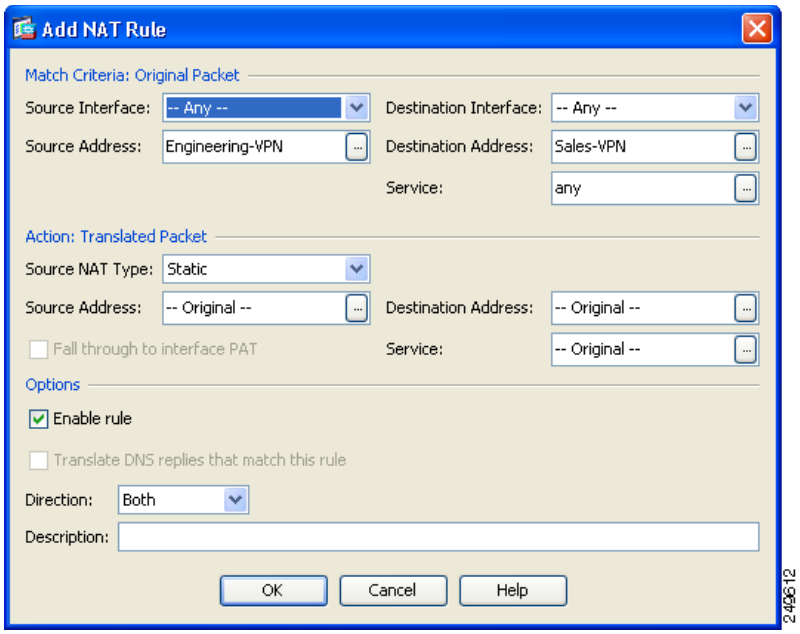

a. In the **Match criteria: Original Packet** area, configure these fields:

- Source Interface: Any
- Destination Interface: Any
- Source Address: Click the Source Address browse button and create the network object that represents the Engineering VPN address pool. Define the object type as a **Range** of addresses. Do not add an automatic address translation rule. See Figure 87-3 for an example.
- Destination Address: Click the Destination Address browse button and create the network object that represents the Sales VPN address pool. Define the object type as a **Range** of addresses. Do not add an automatic address translation rule.

Figure 87-3 Create Network Object for a VPN address pool

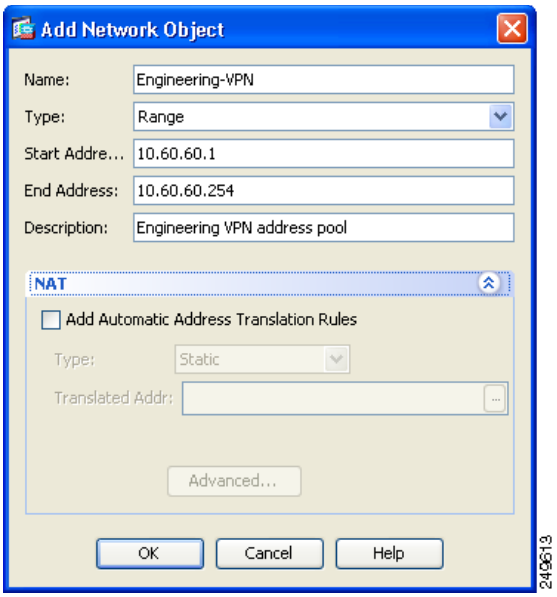

b. In the **Action Translated Packet** area, configure these fields:
   - Source NAT Type: Static
   - Source Address: Original
   - Destination Address: Original
   - Service: Original

c. In the **Options** area, configure these fields:
   - Check **Enable rule**.
   - Uncheck or leave empty the **Translate DNS replies that match this rule**.
   - Direction: Both
   - Description: Add a Description for this rule.

d. Click **OK**.

e. Click **Apply**. Your rule should look like rule 1 in the Unified NAT table in Figure 87-5 on page 87-52.

   CLI example:

```
nat source static Engineering-VPN Engineering-VPN destination static Sales-VPN
Sales-VPN
```

f. Click **Send**.

Step 3 When ASA is performing NAT, in order for two hosts in the same VPN pool to connect to each other, or for those hosts to reach the Internet through the VPN tunnel, you must enable the **Enable traffic between two or more hosts connected to the same interface** option. To do this, in ASDM, select **Configuration > Device Setup > Interfaces**. At the bottom of the Interface panel, check **Enable traffic between two or more hosts connected to the same interface** and click **Apply**.

CLI example:

```
same-security-traffic permit inter-interface
```

Step 4 Create a NAT rule so that the hosts in the Engineering VPN address pool can reach other hosts in the Engineering VPN address pool. Create this rule just as you created the rule in Step 2 except that you specify the Engineering VPN address pool as both the Source address and the Destination Address in the **Match criteria: Original Packet** area.

Step 5 Create a NAT rule so that the Engineering VPN remote access clients can reach the "inside" network. In the NAT Rules pane, select **Add > Add NAT Rule Before "Network Object" NAT rules** so that this rule will be processed before other rules.

a. In the **Match criteria: Original Packet** area configure these fields:

- Source Interface: Any
- Destination Interface: Any
- Source Address: Click the Source Address browse button and create a network object that represents the inside network. Define the object type as a **Network** of addresses. Do not add an automatic address translation rule.
- Destination Address: Click the Destination Address browse button and select the network object that represents the Engineering VPN address pool.

Figure 87-4 Add inside-network object

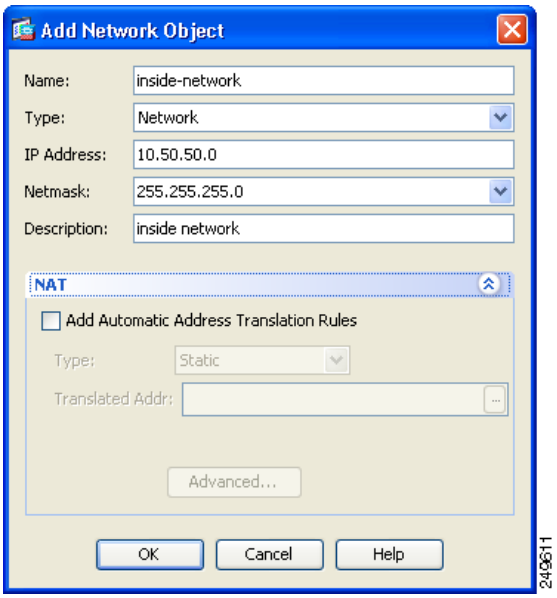

b. In the **Action: Translated Packet** area, configure these fields:

- Source NAT Type: Static
- Source Address: Original
- Destination Address: Original
- Service: Original

c. In the **Options** area, configure these fields:

- Check **Enable rule**.
- Uncheck or leave empty the **Translate DNS replies that match this rule**.
- Direction: Both

- Description: Add a Description for this rule.

d. Click **OK**.

e. Click **Apply**. Your rule should look like rule two in the Unified NAT table in Figure 87-5 on page 87-52.

CLI example

```
nat source static inside-network inside-network destination static Engineering-VPN
Engineering-VPN
```

Step 6 Create a new rule, following the method in Step 5, to configure identity NAT for the connection between the Engineering VPN address pool and the DMZ network. Use the DMZ network as the Source Address and use the Engineering VPN address pool as the Destination address.

Step 7 Create a new NAT rule to allow the Engineering VPN address pool to access the Internet through the tunnel. In this case, you do not want to use identity NAT because you want to change the source address from a private address to an Internet routable address. To create this rule, follow this procedure:

a. In the NAT Rules pane, select **Add > Add NAT Rule Before "Network Object" NAT rules** so that this rule will be processed before other rules.

b. In the **Match criteria: Original Packet** area configure these fields:

- Source Interface: Any
- Destination Interface: Any. This field will be automatically populated with "outside" after you select outside as the Source Address in the **Action: Translated Packet** area.
- Source Address: Click the Source Address browse button and select the network object that represents the Engineering VPN address pool.
- Destination Address: Any.

c. In the **Action: Translated Packet** area, configure these fields:

- Source NAT Type: Dynamic PAT (Hide)
- Source Address: Click the Source Address browse button and select the **outside** interface.
- Destination Address: Original
- Service: Original

d. In the **Options** area, configure these fields:

- Check **Enable rule**.
- Uncheck or leave empty the **Translate DNS replies that match this rule**.
- Direction: Both
- Description: Add a Description for this rule.

e. Click **OK**.

f. Click **Apply**. Your rule should look like rule five in the Unified NAT table in Figure 87-5 on page 87-52.

CLI example:

nat (any,outside) source dynamic Engineering-VPN interface

Figure 87-5 Unified NAT table

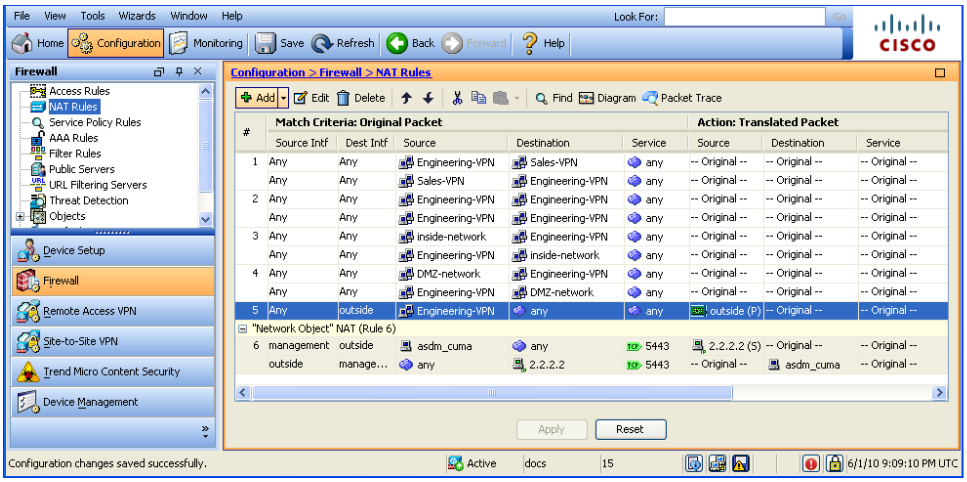

Step 8 After you have configured the Engineering VPN Address pool to reach itself, the Sales VPN address pool, the inside network, the DMZ network, and the Internet; you must repeat this process for the Sales VPN address pool. Use identity NAT to exempt the Sales VPN address pool traffic from undergoing network address translation between itself, the inside network, the DMZ network, and the Internet.

Step 9 From the **File** menu on the ASA, select **Save Running Configuration to Flash** to implement your identity NAT rules.

Configuring AnyConnect VPN Connections

Use the AnyConnect Connection Profiles pane and its child dialog boxes to specify 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—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 specifying whether to allow access.
  - In the Interface table, in the row for the interface you are configuring for AnyConnect connections, check the protocols you want to enable on the interface. You can allow SSL Access, IPsec access, or both.

    When checking SSL, DTLS (Datagram Transport Layer Security) is enabled by default. 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.

When checking IPsec (IKEv2) access, client services are enabled by default. Client services include enhanced Anyconnect features including software updates, client profiles, GUI localization (translation) and customization, Cisco Secure Desktop, and SCEP proxy. If you disable client services, the AnyConnect client still establishes basic IPsec connections with IKEv2.

  - Device Certificate—Lets you specify a certificate for authentication for either an RSA key or an ECDSA key. See Specifying a Device Certificate for AnyConnect Access, page 87-53.
  - Port Setting—Configure port numbers for clientless SSL and IPsec (IKEv2) connections. See Configuring Port Settings, page 87-54.
  - Enable inbound VPN sessions to bypass interface access lists is checked by default.— The security appliance allows all VPN traffic to pass through the interface ACLs. For example, even if the outside interface ACL does not permit the decrypted traffic to pass through, the security appliance trusts the remote private network and permits the decrypted packets to pass through. You can change this default behavior. If you want the interface ACL to inspect the VPN protected traffic, uncheck this box.

- 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.
  - Shutdown portal login page.—Shows the web page when the login is disabled.
- 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.
- Let group URL take precedence if group URL and certificate map match different connection profiles. Otherwise, the connection profile matches the certificate map will be used.—This option specifies the relative preference of the group URL and certificate values during the connection profile selection process. If the ASA fails to match the preferred value, it chooses the connection profile that matches the other value. Check this option only if you want to rely on the preference used by many older ASA software releases to match the group URL specified by the VPN endpoint to the connection profile that specifies the same group URL. This option is unchecked by default. If it is unchecked, the ASA prefers to match the certificate field value specified in the connection profile to the field value of the certificate used by the endpoint to assign the connection profile.

Specifying a Device Certificate for AnyConnect Access

The **Specify Device Certificate** screen allows you to specify an RSA or ECDSA certificate that will identify the ASA to the AnyConnect client when it attempts to create a VPN connection. Certain AnyConnect features, such as Always-on IPsec/IKEv2, require that a valid and trusted device certificate be available on the ASA.

If your AnyConnect clients are configured to use only SSL, then you only need to specify an RSA certificate as AnyConnect does not support ECDSA certificates for SSL VPN. If your AnyConnect clients are configured to use IPsec or SSL, or both IPsec and SSL, you can configure both kinds of certificates.

ECDSA certificates are only supported on IPsec connections.

Prerequisites

RSA or ECDSA certificates need to be stored on the ASA.

Detailed Steps

Step 1 In the Certificate with RSA Key area, perform one of these tasks:

- Keep the **Use the same device certificate for SSL and IPsec IKEv2** box checked if you want to choose one certificate to authenticate clients using either protocol. You can select the certificate from those available in the list box or click **Manage** to create an identity certificate to use.
- Uncheck the **Use the same device certificate for SSL and IPsec IKEv2** check box to specify separate certificates for SSL connections or IPsec connections.

Step 2 In the Certificate with ECDSA key field select the ECDSA certificate from the list box or click **Manage** to create an ECDSA identity certificate.

Step 3 Click **OK**.

Configuring Port Settings

Configure port numbers for SSL and IPsec (IKEv2) connections in this window:

Fields

- SSL Ports:
  - HTTPS Port—The port to enable HTTPS for clientless (browser-based) SSL connections. The range is 1-65535. The default is port 443.
  - DTLS Port—The port to enable DTLS for SSL connections. The range is 1-65535. The default is port 443.
- IPsec Client Services Port—The port to enable client services for IKEv2 connections. The range is 1-65535. The default is port 443.

Setting the Basic Attributes for an AnyConnect VPN Connection

To set the basic attributes for an AnyConnect VPN connection, choose Add or Edit in the Anyconnect Connection Profiles section. The Add (or Edit) Anyconnect Connection Profile > Basic dialog box opens.

Fields

Set the attributes in the Add AnyConnect Connection Profile > 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 ASA 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 ASA 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 pool name of an available, configured pool of IPv4 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. See Configuring Local IP Address Pools, page 88-3 for more information on adding or editing an IPv4 address pool.
  - Client IPv6 Address Pools—Enter the pool name of an available, configured pool of IPv6 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. See Configuring Local IP Address Pools, page 88-3 for more information on adding or editing an IPv6 address pool.
- 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 to enable SSL for this VPN connection.
  - Enable IPsec (IKEv2) client protocol—Check to enable IPsec using IKEv2 for this connection.
  - DNS Servers—Enter the IP address(s) of DNS servers for this policy.
  - WINS Servers—Enter the IP address(s) of WINS servers for this policy.
  - Domain Name—Enter a default domain name.
- Find—Enter a GUI label or a CLI command to use as a search string, then click Next or Previous to begin the search.

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 Simple Certificate Enrollment (SCEP) for this Connection Profile
- Strip the realm from username before passing it on to the AAA server
- Strip the group from username before passing it on to the AAA server
- Group Delimiter—Changing the group delimiter value makes the change globally on all other remote connection profiles.
- 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 ASA 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 ASA 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 ASA, 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 ASA 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 ASA 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 ASA 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.

- Translate Assigned IP Address to Public IP Address—In rare situations, you might want to use a VPN peer's real IP address on the inside network instead of an assigned local IP address. Normally with VPN, the peer is given an assigned local IP address to access the inside network. However, you might want to translate the local IP address back to the peer's real public IP address if, for example, your inside servers and network security is based on the peer's real IP address. You can enable this feature on one interface per tunnel group.
  - Enable the address translation on interface—Enables the address translation and allows you to choose which interface the address appears on. *Outside* is the interface to which the AnyConnect client connects, and *inside* is the interface specific to the new tunnel group.

Note Because of routing issues and other limitations, we do not recommend using this feature unless you know you need 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.

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 87-87.

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.

Attribute	Definition
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.

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 ASA 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.
- Username Mapping from Certificate—Specify the fields in a digital certificate from which to extract the username.
- Pre-fill Username from Certificate—Check to extract the names to be used for secondary authentication from the primary and secondary fields specified in this panel. You must configure the authentication method for both AAA and certificates before checking this attribute. To do so, return to the Basic panel in the same window and check Both next to Method.
- Hide username from end user—Check to hide the username to be used for secondary authentication from the VPN user.
- Fallback when a certificate is unavailable —This attribute is configurable only if “Hide username from end user” is checked. Uses Cisco Secure Desktop Host Scan data to pre-fill the username for secondary authentication if a certificate is unavailable.
- Password—Choose one of the following methods to retrieve the password to be used for secondary authentication:
  - Prompt—Prompt the user for the password.
  - Use Primary—Reuse the primary authentication password for all secondary authentications.
  - Use—Enter a common secondary password for all secondary authentications.
- 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.
  - 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.

  The options for primary and secondary field 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.

- Use the entire DN as the username—Uses the entire subject DN (RFC1779) to derive a name for an authorization query from a digital certificate.
- Use script to select username—Names the script from which to extract 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.

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.

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 72-1 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 87-3 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@example.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 87-4 lists the attribute names and descriptions that you can use in a Lua script.

Note Lua is case-sensitive.

Table 87-4 Attribute Names and Descriptions

Attribute Name	Description
cert.subject.c	Country
cert.subject.cn	Common Name
cert.subject.dnq	DN qualifier
cert.subject.ea	E-mail Address
cert.subject.genq	Generational qualified

Table 87-4 Attribute Names and Descriptions

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	E-mail Address
cert.issuer.genq	Generational qualified
cert.issuer.gn	Given Name
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.

Configuring AnyConnect Secure Mobility

AnyConnect Secure Mobility protects corporate interests and assets from Internet threats when employees are mobile. Use the Mobile User Security dialog box to configure this feature. AnyConnect Secure Mobility lets Cisco IronPort S-Series Web Security appliances scan Cisco AnyConnect secure

mobility clients to ensure that clients are protected from malicious software and/or inappropriate sites. The client periodically checks to ensure that Cisco IronPort S-Series Web Security appliance protection is enabled.

To configure secure mobility solutions, choose **Configuration > Remote Access VPN > Network (Client) Access > Mobile User Security**.

Note This feature requires a release of the Cisco IronPort Web Security appliance that provides AnyConnect Secure Mobility licensing support for the Cisco AnyConnect secure mobility client. It also requires an AnyConnect release that supports the AnyConnect Secure Mobility feature.

Figure 87-6 Mobile User Security Window

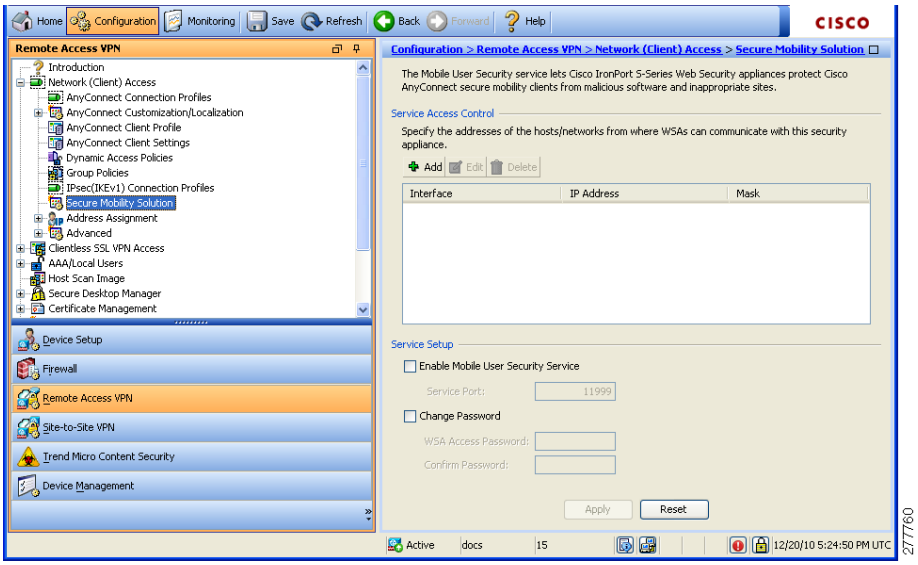

Fields

- Service Access Control—Specifies from which host or network address the WSAs can communicate.
  - Add—Opens the Add MUS Access Control Configuration dialog box for the selected connection.
  - Edit—Opens the Edit MUS Access Control Configuration dialog box for the selected connection.
  - Delete—Removes the selected connection from the table. There is no confirmation or undo.
- Enable Mobile User Security Service—Starts the connection with the client through the VPN. If enabled, you are required to enter a password, used by the WSA when contacting the ASA. If no WSA is present, the status is disabled.
- Service Port—If you choose to enable the service, specify which port number for the service to use. The port must be between 1 and 65535 and must match the corresponding value provisioned into the WSA with the management system. The default is 11999.
- Change Password—Enables you to change the WSA access password.

- WSA Access Password—Specify the shared secret password required for authentication between the ASA and WSA. This password must match the corresponding password provisioned into the WSA with the management system.
- Confirm Password—Re-enter the specified password.
- Show WSA Sessions—Allows you to view session information of WSAs connected to the ASA.The host IP address of the WSA that is connected (or has been connected) and the duration of the connection is returned in a dialog box.

Add or Edit MUS Access Control

The Add or Edit MUS Access Control dialog box lets you configure MUS access.

Fields

- Interface Name—Use the drop-down menu to choose which interface name you are adding or editing.
- IP Address—Enter either an IPv4 or IPv6 address.
- Mask—Use the drop-down menu to choose the appropriate mask.

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 specifying whether to allow access.
  - Device Certificate—Lets you specify a certificate for authentication for either an RSA key or an ECDSA key or trustpoint. You have the option to configure two trustpoints. The client indicates ECDSA support with a vendor ID payload. The ASA scans the configured trustpoint list and chooses the first one that the client supports. If ECDSA is preferred, you should configure that trustpoint before the RSA trustpoint.
  - Manage—Opens the Manage Identity Certificates dialog box, on which you can add, edit, delete, export, and show details for a selected certificate.
  - Port Setting—Configure port numbers for clientless SSL and IPsec (IKEv2) connections. The range is 1-65535. The default is port 443.
- Login Page Setting
  - Allow user to select connection profile, identified by its alias, on the login page. Otherwise, DefaultWebVPN Group will be the connection profile.—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.
  - Allow user to enter internal password on the login page.—Adds an option to input a different password when accessing internal servers.
  - Shutdown portal login page.—Shows the web page when the login is disabled.

- Connection Profiles—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.
  - Name—The name of the Connection Profile.
  - Enabled—Checkmark when enabled.
  - Aliases—Other names by which the Connection Profile is known.
  - Authentication Method—Specifies which authentication method is used.
  - Group Policy—Shows the default group policy for this Connection Profile.
- Let group URL take precedence if group URL and certificate map match different connection profiles. Otherwise, the connection profile matches the certificate map will be used.—This option specifies the relative preference of the group URL and certificate values during the connection profile selection process. If the ASA fails to match the preferred value specified by the endpoint to that specified by a connection profile, it chooses the connection profile that matches the other value. Check this option only if you want to rely on the preference used by many older ASA software releases to match the group URL specified by the VPN endpoint to the connection profile that specifies the same group URL. This option is unchecked by default. If it is unchecked, the ASA prefers to match the certificate field value specified in the connection profile to the field value of the certificate used by the endpoint to assign the connection profile.

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.

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 ASA 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 ASA 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.

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.

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.

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.

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.

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.

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.

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 ASA 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 ASA 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 ASA 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.

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.

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.

IPsec Remote Access Connection Profiles

Configuration > VPN > General > Tunnel Group

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.

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.
  - 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.

Mapping Certificates to IPsec or SSL VPN Connection Profiles

When the ASA 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 ASA only uses the rules you configure.

For IPsec or SSL connections using rules, the ASA 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 ASA 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

Configuration > VPN > IKE > Certificate Group Matching > Rules > 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 ASA evaluates the map when it receives a connection request. For the first rule defined, the default priority is 10. The ASA 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 ASA ignores the map entry.

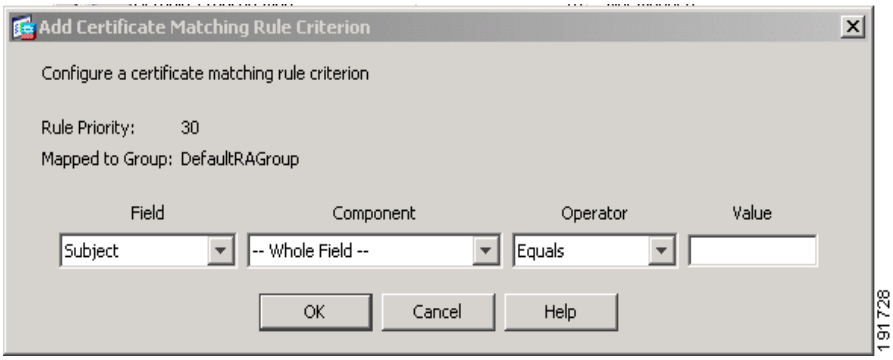

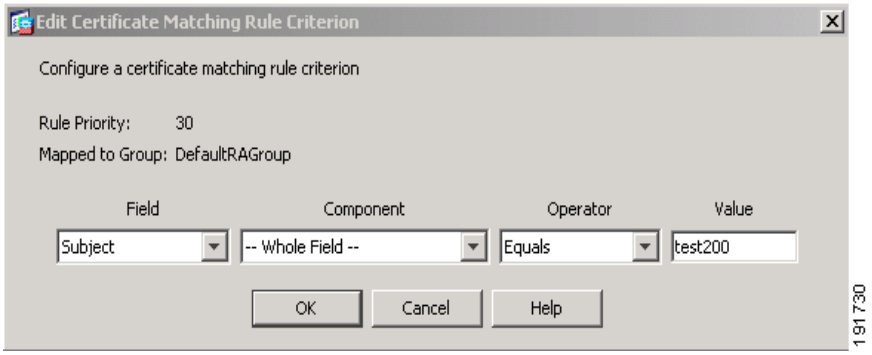

Add/Edit Certificate Matching Rule Criterion

Configuration > VPN > IKE > Certificate Group Matching > Rules > Add/Edit Certificate Matching Rule Criterion

Use the **Add/Edit Certificate Matching Rule Criterion** dialog box to configure a certificate matching rule criterion for the selected connection profile.

Fields

- **Rule Priority**—(Display only). Sequence with which the ASA evaluates the map when it receives a connection request. The ASA 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
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

Site-to-Site Connection Profiles

The Connection Profiles dialog box shows the attributes of the currently configured Site-to-Site connection profiles (tunnel groups), lets you select the delimiter to use when parsing connection profile names, and lets you add, modify, or delete connection profiles.

The security appliance supports IPsec LAN-to-LAN VPN connections for IPv4 or IPv6 using IKEv1 or IKEv2 and supports both inside and outside networks using the inner and outer IP headers.

Fields

- Access Interfaces—Displays a table of device interfaces where you can enable access by a remote peer device on the interface:
  - Interface—The device interface to enable or disable access.
  - Allow IKEv1 Access—Check to enable IPsec IKEv1 access by a peer device.
  - Allow IKEv2 Access—Check to enable IPsec IKEv2 access by a peer device.
- Connection Profiles—Displays a table of connection profiles where you can add, edit, or delete profiles:
  - Add—Opens the Add IPsec Site-to-Site connection profile dialog box.
  - Edit—Opens the Edit IPsec Site-to-Site connection profile dialog box.
  - Delete—Removes the selected connection profile. There is no confirmation or undo.
  - Name—The name of the connection profile.
  - Interface—The interface the connection profile is enabled on.
  - Local Network—Specifies the IP address of the local network.
  - Remote Network—Specifies the IP address of the remote network.
  - IKEv1 Enabled—Shows IKEv1 enabled for the connection profile.
  - IKEv2 Enabled—Shows IKEv2 enabled for the connection profile.
  - Group Policy—Shows the default group policy of the connection profile.

Add/Edit Site-to-Site Connection

You can get to this panel through various paths.

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 (IPv4 or IPv6), specify a connection name, select an interface, specify IKEv1 and IKEv2 peer and user authentication parameters, specify protected networks, and specify encryption algorithms.

The ASA supports LAN-to-LAN VPN connections to Cisco or third-party peers when the two peers have IPv4 inside and outside networks (IPv4 addresses on the inside and outside interfaces).

For LAN-to-LAN connections using mixed IPv4 and IPv6 addressing, or all IPv6 addressing, the security appliance supports VPN tunnels if both peers are Cisco ASA 5500 series security appliances, and if both inside networks have matching addressing schemes (both IPv4 or both IPv6).

Specifically, the following topologies are supported when both peers are Cisco ASA 5500 series ASAs:

- The ASAs have IPv4 inside networks and the outside network is IPv6 (IPv4 addresses on the inside interfaces and IPv6 addresses on the outside interfaces).
- The ASAs have IPv6 inside networks and the outside network is IPv4 (IPv6 addresses on the inside interface and IPv4 addresses on the outside interfaces).
- The ASAs have IPv6 inside networks and the outside network is IPv6 (IPv6 addresses on the inside and outside interfaces).

Fields

- Peer IP Address—Lets you specify an IP address (IPv4 or IPv6) and whether that address is static.
- Connection Name—Specifies the name assigned to this connection profile. 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.
- Protected Networks—Selects or specifies the local and remote network protected for this connection.
  - IP Address Type—Specifies the address is an IPv4 or IPv6 address.
  - 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.
- IPsec Enabling—Specifies the group policy for this connection profile and the key exchange protocol specified in that policy:
  - Group Policy Name—Specifies the group policy associated with this connection profile.
  - Manage—Opens the Browse Remote Network dialog box, in which you can select a remote network.
  - Enable IKEv1—Enables the key exchange protocol IKEv1 in the specified group policy.
  - Enable IKEv2—Enables the key exchange protocol IKEv2 in the specified group policy.
- IKEv1 Settings tab—Specifies authentication and encryption settings for IKEv1:
  - 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.

  - Device Certificate—Specifies the name of the identity certificate, if available, to use for authentication.
  - 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 Policy—Specifies one or more encryption algorithms to use for the IKE proposal.
  - Manage—Opens the Configure IKEv1 Proposals dialog box.
  - IPsec Proposal—Specifies one or more encryption algorithms to use for the IPsec IKEv1 proposal.
- IKEv2 Settings tab—Specifies authentication and encryption settings for IKEv2:
  - Local 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.
  - Local Device Certificate—Specifies the name of the identity certificate, if available, to use for authentication.
  - 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.
  - Remote Peer Pre-shared Key—Specify the value of the remote peer pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.
  - Remote Peer Certificate Authentication—Check *Allowed* to allow certificate authentication for IKEv2 connections for this connection profile.
  - Manage—Opens the Manage CA Certificates dialog where you can view certificates and add new ones.
  - IKE Policy—Specifies one or more encryption algorithms to use for the IKE proposal.
  - Manage—Opens the Configure IKEv1 Proposals dialog box.
  - IPsec Proposal—Specifies one or more encryption algorithms to use for the IPsec IKEv1 proposal.
  - Select—Opens the Select IPsec Proposals (Transform Sets) dialog box, where you can assign a proposal to the connection profile for IKEv2 connections.

Adding or Editing a Site-to-Site Tunnel Group

You can get to this panel through various paths.

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.
- IPsec Enabling—Specifies the group policy for this connection profile and the key exchange protocol specified in that policy:
  - Group Policy Name—Specifies the group policy associated with this connection profile.
  - Manage—Opens the Browse Remote Network dialog box, in which you can select a remote network.
  - Enable IKEv1—Enables the key exchange protocol IKEv1 in the specified group policy.
  - Enable IKEv2—Enables the key exchange protocol IKEv2 in the specified group policy.
- IKEv1 Settings tab—Specifies authentication and encryption settings for IKEv1:
  - 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.
  - Device Certificate—Specifies the name of the identity certificate, if available, to use for authentication.
  - 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 Policy—Specifies one or more encryption algorithms to use for the IKE proposal.
  - Manage—Opens the Configure IKEv1 Proposals dialog box.
  - IPsec Proposal—Specifies one or more encryption algorithms to use for the IPsec IKEv1 proposal.
- IKEv2 Settings tab—Specifies authentication and encryption settings for IKEv2:
  - Local 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.
  - Local Device Certificate—Specifies the name of the identity certificate, if available, to use for authentication.
  - 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.
  - Remote Peer Pre-shared Key—Specify the value of the remote peer pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.
  - Remote Peer Certificate Authentication—Check *Allowed* to allow certificate authentication for IKEv2 connections for this connection profile.
  - Manage—Opens the Manage CA Certificates dialog where you can view certificates and add new ones.
  - IKE Policy—Specifies one or more encryption algorithms to use for the IKE proposal.
  - Manage—Opens the Configure IKEv1 Proposals dialog box.

  - IPsec Proposal—Specifies one or more encryption algorithms to use for the IPsec IKEv1 proposal.
  - Select—Opens the Select IPsec Proposals (Transform Sets) dialog box, where you can assign a proposal to the connection profile for IKEv2 connections.
- 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 ASA 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 ASA never initiates keepalive monitoring.

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.

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.

:

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@example.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 ASA 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@example.com#VPNGroup*. If you choose this option, you must use either the # or ! character for the group delimiter because the ASA 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 example.com domain, you might call your Kerberos realm EXAMPLE.COM.

The ASA 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 ASA 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 ASA 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 ASA ignores this command if RADIUS or LDAP authentication has not been configured.

This feature requires the use of MS-CHAPv2.

Configuring Client Addressing

To specify the client IP address assignment policy and assign address pools to all IPsec and SSL VPN connections, open ASDM and select Configuration > 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 ASA. 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 ASA, choose **Add**. The Add IP Pool dialog box opens.
  – To change the configuration of an address pool on the ASA, 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 ASA, 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 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 ASA uses the selected sources in order, until it finds an address:
  - Use authentication server—Specifies that the ASA should attempt to use the authentication server as the source for a client address.
  - Use DHCP—Specifies that the ASA should attempt to use DHCP as the source for a client address.

  - Use address pool—Specifies that the ASA should attempt to use address pools as the source for a client address.
- Interface-Specific Address Pools—Lists the configured interface-specific address pools.

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.

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.

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.

Add/Edit Connection Profile > General > Authentication

You can get to this panel through various paths.

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 ASA. 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.

Add/Edit SSL VPN Connection > General > Authorization

You can get to this panel through various paths.

The settings on this dialog box apply to the connection (tunnel group) globally across the ASA. 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 ASA 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 ASA 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

You can get to this panel through various paths.

The settings on this dialog box apply to the connection (tunnel group) globally across the ASA. 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.

Add/Edit Tunnel Group > General > Client Address Assignment

You can get to this panel through various paths.

To specify whether to use DHCP or address pools for address assignment, go to Configuration > VPN > I P Address Management > Assignment. The Add or Edit Tunnel Group 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.

Add/Edit Tunnel Group > General > Advanced

You can get to this panel through various paths.

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.

Add/Edit Tunnel Group > IPsec for Remote Access > IPsec

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for Remote Access > IPSec Tab

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 ASA 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 ASA 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://.

Add/Edit Tunnel Group for Site-to-Site VPN

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for Remote Access > IPSec Tab

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 ASA 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 ASA 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.

Add/Edit Tunnel Group > PPP

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > PPP Tab

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 MS-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

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for LAN to LAN Access > General Tab > Basic Tab

On the Add or Edit Tunnel Group 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@example.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 ASA 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@example.com#VPNGroup*. If you choose this option, you must use either the # or ! character for the group delimiter because the ASA 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 example.com domain, you might call your Kerberos realm EXAMPLE.COM.

The ASA 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 ASA 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 ASA 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 ASA 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.

Add/Edit Tunnel Group > IPsec for LAN to LAN Access > IPsec

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for LAN to LAN Access > IPSec Tab

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 ASA 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 ASA 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://.

Clientless SSL VPN Access > Connection Profiles > Add/Edit > General > Basic

Configuration > VPN > General > Tunnel Group > Add/Edit> WebVPN Access > General Tab > Basic Tab

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 ASA 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 ASA 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.

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@example.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 ASA 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@example.com#VPNGroup*. If you choose this option, you must use either the # or ! character for the group delimiter because the ASA 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 example.com domain, you might call your Kerberos realm EXAMPLE.COM.

The ASA 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 ASA 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 ASA 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 ASA ignores this command if RADIUS or LDAP authentication has not been configured.

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

- Interface-Specific IPv4 Address Pools—Lists the configured interface-specific address pools.
- Interface-Specific IPv6 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.

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.

Authenticating SSL VPN Connections

The SSL VPN Connections > Advanced > Authentication dialog box lets you configure authentication attributes for SSL VPN connections.

System Options

This field can be reached by navigating these paths:

- Configuration > Site-to-Site VPN > Advanced > System Options
- Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > System Options

The System Options pane lets you configure features specific to VPN sessions on the ASA.

Fields

- 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 IPsec 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. This is an advanced system option for Network (Client) Access only.
- Reclassify existing flows when VPN tunnels establish
- Preserve stateful VPN flows when the tunnel drops—Enables or disables preserving IPsec tunneled flows in Network-Extension Mode (NEM). 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).

- IPsec 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.
- Enable PMTU (Path Maximum Transmission Unit) Aging—Allows an administrator to enable PMTU aging.
  - Interval to Reset PMTU of an SA (Security Association)—Enter the number of seconds at which the PMTU value is reset to its original value.

Zone Labs Integrity Server

Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > Zone Labs Integrity Server

The Zone Labs Integrity Server panel lets you configure the ASA 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 ASA 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.

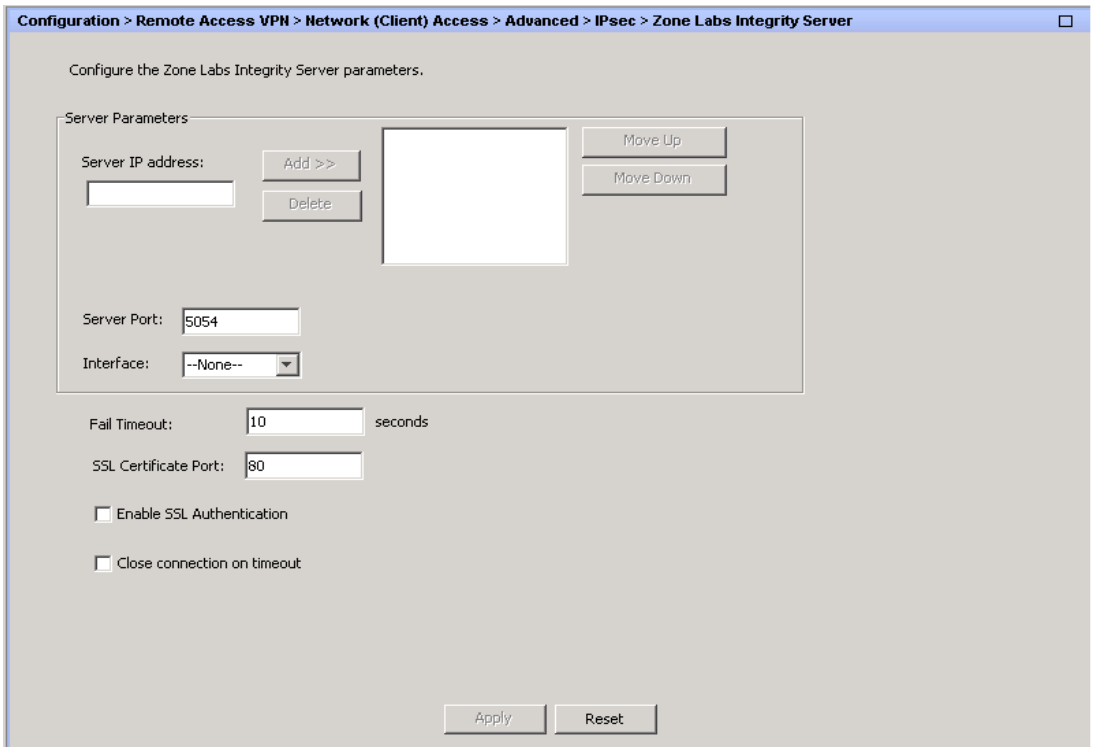

Note The current release of the security appliance supports one Integrity Server at a time even though the user interfaces support the configuration of up to five Integrity Servers. If the active Server fails, configure another Integrity Server on the ASA and then reestablish the client VPN session.

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 ASA 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 ASA 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 ASA 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 ASA 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 ASA. By default, client SSL authentication is disabled.
- Close connection on timeout—Check to close the connection between the ASA and the Integrity Server on a timeout. By default, the connection remains open.
- Apply—Click to apply the Integrity Server setting to the ASA running configuration.
- Reset—Click to remove Integrity Server configuration changes that have not yet been applied.

Easy VPN Remote

Configuration > VPN > Easy VPN Remote

Easy VPN Remote lets the ASA 5505 act as an Easy VPN client device. The ASA 5505 can then initiate a VPN tunnel to an Easy VPN server, which can be an ASA, a Cisco VPN 3000 Concentrator, a Cisco 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 ASA 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 ASA 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.

Advanced Easy VPN Properties

Configuration > VPN > Easy VPN Remote > Advanced

Device Pass-Through

Certain devices like Cisco IP phones, printers, and the like are incapable of performing authentication, and therefore of participating in individual unit authentication. To accommodate these devices, the device pass-through feature, enabled by the MAC Exemption attributes, exempts devices with the specified MAC addresses from authentication when Individual User Authentication is enabled.

The first 24 bits of the MAC address indicate the manufacturer of the piece of equipment. The last 24 bits are the unit's serial number in hexadecimal format.

Tunneled Management

When operating an ASA model 5505 device behind a NAT device, use the Tunneled Management attributes to specify how to configure device management— in the clear or through the tunnel—and specify the network or networks allowed to manage the Easy VPN Remote connection through the tunnel. The public address of the ASA 5505 is not accessible when behind the NAT device unless you add static NAT mappings on the NAT device.

When operating a Cisco ASA 5505 behind a NAT device, use the **vpnclient management** command to specify how to configure device management— with additional encryption or without it—and specify the hosts or networks to be granted administrative access. The public address of the ASA 5505 is not accessible when behind the NAT device unless you add static NAT mappings on the NAT device.

Fields

- MAC Exemption—Configures a set of MAC addresses and masks used for device pass-through for the Easy VPN Remote connection
  - MAC Address—Exempts the device with the specified MAC address from authentication. The format for specifying the MAC address this field uses three hex digits, separated by periods; for example, 45ab.ff36.9999.
  - MAC Mask—The format for specifying the MAC mask in this field uses three hex digits, separated by periods; for example, the MAC mask ffff.ffff.ffff matches just the specified MAC address. A MAC mask of all zeroes matches no MAC address, and a MAC mask of ffff.ff00.0000 matches all devices made by the same manufacturer.
  - Add—Adds the specified MAC address and mask pair to the MAC Address/Mask list.
  - Remove—Moves the selected MAC address and mask pair from the MAC Address/MAC list to the individual MAC Address and MAC Mask fields.
- Tunneled Management—Configures IPsec encryption for device management and specifies the network or networks allowed to manage the Easy VPN hardware client connection through the tunnel. Selecting Clear Tunneled Management merely removes that IPsec encryption level and does not affect any other encryption, such as SSH or https, that exists on the connection.
  - Enable Tunneled Management—Adds a layer of IPsec encryption to the SSH or HTTPS encryption already present in the management tunnel.
  - Clear Tunneled Management—Uses the encryption already present in the management tunnel, without additional encryption.
  - IP Address— Specifies the IP address of the host or network to which you want to grant administrative access to the Easy VPN hardware client through the VPN tunnel. You can individually add one or more IP addresses and their respective network masks.

  - Mask—Specifies the network mask for the corresponding IP address.
  - Add—Moves the specified IP address and mask to the IP Address/Mask list.
  - Remove—Moves the selected IP address and mask pair from the IP Address/Mask list to the individual IP Address and Mask fields in this area.
  - IP Address/Mask—Lists the configured IP address and mask pairs to be operated on by the Enable or Clear functions in this area.
- IPsec Over TCP—Configure the Easy VPN Remote connection to use TCP-encapsulated IPsec.
  - Enable—Enables IPsec over TCP.

Note Choose Configuration > VPN > IPsec > Pre-Fragmentation, double-click the outside interface, and set the DF Bit Setting Policy to Clear if you configure the Easy VPN Remote connection to use TCP-encapsulated IPsec. The Clear setting lets the ASA 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.

AnyConnect Custom Attributes

Custom attributes are added here to support special features that are not defined in the ASDM. Deferred upgrade is the feature in AnyConnect 3.1 that uses custom attributes. See Configuring AnyConnect Client Custom Attributes for an Internal Group Policy, page 87-17

AnyConnect Essentials

AnyConnect Essentials is a separately licensed SSL VPN client, entirely configured on the ASA, 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 ASA.

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.

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 ASA.
- 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.

SSL VPN Client Settings

The Cisco AnyConnect VPN client provides secure SSL connections to the ASA 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 ASA is configured to redirect http:// requests to https://, users must enter the URL in the form https://<*address*>.

If you need to redirect http:// requests to https://, go to **Configuration > Remote Access VPN > Advanced**, then click **HTTP Redirect**. Choose the interface you want to redirect, then click **Edit** to display the Edit HTTP/HTTPS Settings dialog box. Check the **Redirect HTTP to HTTPS** check box, and change the HTTP Port value, if necessary. Click **OK** to close this dialog box, then click **Apply** to save your settings.

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 ASA 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 ASA configuration) when the connection terminates.

In the case of a previously installed client, when the user authenticates, the ASA examines the revision of the client, and upgrades the client as necessary.

When the client negotiates an SSL VPN connection with the ASA, 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 ASA, 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 ASA downloads the client based on the group policy or local user policy attributes. You can configure the ASA 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 ASA 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 ASA 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 highlighted 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 ASA 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.

Add/Replace SSL VPN Client Image

In this pane, you can specify a filename for a file on the ASA 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 ASA 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 ASA, it includes the User-Agent string in the HTTP header. When the ASA receives the string, if the string matches an expression configured for an image, it immediately downloads that image without testing the other client images.

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.

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 ASA 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.

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 ASA allows VPN traffic to terminate on an ASA 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 ASA performance without any security risks. (Group policy and per-user authorization access lists still apply to the traffic.)

Configuring AnyConnect Host Scan

Configuration > Remote Access VPN > Host Scan Image

The AnyConnect Posture Module provides the AnyConnect Secure Mobility Client the ability to identify the operating system, anti-virus, anti-spyware, and firewall software installed on the host. The Host Scan application gathers this information.

Using the secure desktop manager tool in the Adaptive Security Device Manager (ASDM), you can create a prelogin policy which evaluates the operating system, anti-virus, anti-spyware, and firewall software Host Scan identifies. Based on the result of the prelogin policy's evaluation, you can control which hosts are allowed to create a remote access connection to the security appliance.

The Host Scan support chart contains the product name and version information for the anti-virus, anti-spyware, and firewall applications you use in your prelogin policies. We deliver Host Scan and the Host Scan support chart, as well as other components, in the Host Scan package.

Starting with AnyConnect Secure Mobility Client, release 3.0, Host Scan is available separately from CSD. This means you can deploy Host Scan functionality without having to install CSD and you will be able to update your Host Scan support charts by upgrading the latest Host Scan package.

Posture assessment and the AnyConnect telemetry module require Host Scan to be installed on the host.

This chapter contains the following sections:

- Host Scan Dependencies and System Requirements, page 87-113
- Host Scan Packaging, page 87-114
- Installing and Enabling Host Scan on the ASA, page 87-115
- Other Important Documentation Addressing Host Scan, page 87-119

Host Scan Dependencies and System Requirements

Dependencies

The AnyConnect Secure Mobility Client with the posture module requires these minimum ASA components:

- ASA 8.4
- ASDM 6.4

These AnyConnect features require that you install the posture module.

- SCEP authentication
- AnyConnect Telemetry Module

System Requirements

The posture module can be installed on any of these platforms:

- Windows XP (x86 and x86 running on x64)
- Windows Vista (x86 and x86 running on x64)
- Windows 7 (x86 and x86 running on x64)
- Mac OS X 10.5,10.6 (32-bit and 32-bit running on 64-bit)
- Linux (32-bit and 32-bit running on 64-bit)
- Windows Mobile

Licensing

These are the AnyConnect licensing requirements for the posture module:

- AnyConnect Premium for basic Host Scan.
- Advanced Endpoint Assessment license is required for
  - Remediation
  - Mobile Device Management

Entering an Activation Key to Support Advanced Endpoint Assessment

Advanced Endpoint Assessment includes all of the Endpoint Assessment features and lets you configure an attempt to update noncompliant computers to meet version requirements. You can use ASDM to activate a key to support Advanced Endpoint Assessment after acquiring it from Cisco, as follows:

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

Step 2 Enter the key in the **New Activation Key** field.

Step 3 Click **Update Activation Key**.

Step 4 Choose **File > Save Running Configuration to Flash**.

An Advanced Endpoint Assessment entry appears and the Configure button becomes active in the Host Scan Extensions area of the **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan** pane, which is accessible only if CSD is enabled.

Host Scan Packaging

You can load the Host Scan package on to the ASA in one of these ways:

- You can upload it as a standalone package: **hostscan-*version*.pkg**
- You can upload it by uploading an AnyConnect Secure Mobility package: **anyconnect-NGC-win-*version*-k9.pkg**
- You can upload it by uploading a Cisco Secure Desktop package: **csd_version-k9.pkg**

File	Description
hostscan-*version*.pkg	This file contains the Host Scan software as well as the Host Scan library and support charts.
anyconnect-NGC-win-*version*-k9.pkg	This package contains all the Cisco AnyConnect Secure Mobility Client features including the hostscan-*version*.pkg file.
csd_*version*-k9.pkg	This file contains all Cisco Secure Desktop features including Host Scan software as well as the Host Scan library and support charts. This method requires a separate license for Cisco Secure Desktop.

Installing and Enabling Host Scan on the ASA

These tasks describe installing and enabling Host Scan on the ASA:

- Installing or Upgrading Host Scan
- Enabling or Disabling Host Scan
- Enabling or Disabling CSD on the ASA
- Viewing the Host Scan Version Enabled on the ASA
- Uninstalling Host Scan
- Uninstalling CSD from the ASA
- Assigning AnyConnect Posture Module to a Group Policy

Installing or Upgrading Host Scan

Use this procedure to upload, or upgrade, and enable a new Host Scan image on the ASA. This image can enable the host scan functionality for AnyConnect, or you can use it to upgrade the host scan support chart for an existing deployment of Cisco Secure Desktop (CSD).

You can specify a standalone Host Scan package or an AnyConnect Secure Mobility Client version 3.0 or later package in the field.

If you previously uploaded a CSD image to the ASA, the Host Scan image you specify will upgrade or downgrade the existing Host Scan files that were delivered with that CSD package.

You do not need to restart the security appliance after you install or upgrade Host Scan; however, you must exit and restart Adaptive Security Device Manager (ASDM) to access Secure Desktop Manager.

Note Host scan requires an AnyConnect Secure Mobility Client premium license.

Step 1 Use your Internet browser to download the **hostscan_*version*-k9.pkg** file or **anyconnect-NGC-win-*version*-k9.pkg** file to your computer. You cannot use a csd_*version*-k9.pkg with this procedure.

Step 2 Open ASDM and choose **Configuration > Remote Access VPN > Host Scan Image.** ASDM opens the Host Scan Image panel (Figure 87-7).

Figure 87-7 Host Scan Image Panel

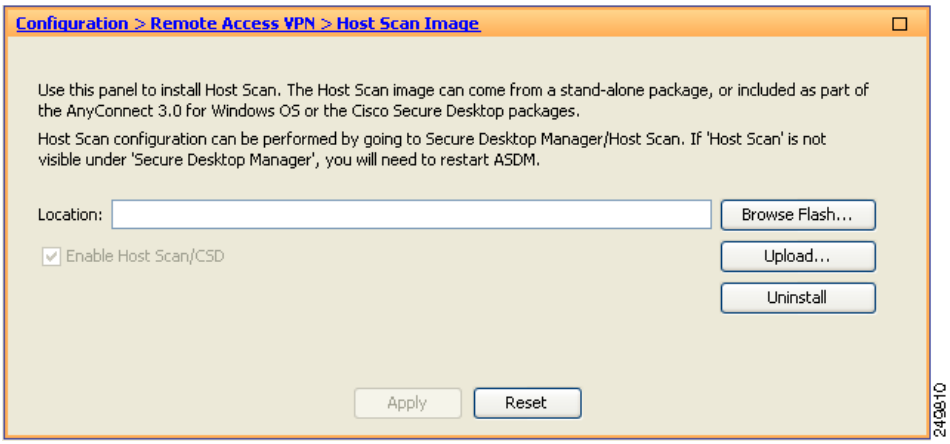

Step 3 Click **Upload** to prepare to transfer a copy of the Host Scan package from your computer to a drive on the ASA.

Step 4 In the Upload Image dialog box, click **Browse Local Files** to search for the Host Scan package on your local computer.

Step 5 Select the **hostscan_*version*.pkg** file or **anyconnect-NGC-win-*version*-k9.pkg** file you downloaded in Step 1 and click **Select**. The path to the file you selected is in the Local File Path field and the Flash File System Path field reflects the destination path of the Host Scan package. If your ASA has more than one flash drive, you can edit the Flash File System Path to indicate another flash drive.

Step 6 Click **Upload File**. ASDM transfers a copy of the file to the flash card. An Information dialog box displays the following message:

```
File has been uploaded to flash successfully.
```

Step 7 Click **OK**.

Step 8 In the Use Uploaded Image dialog, click **OK** to use the Host Scan package file you just uploaded as the current image.

Step 9 Check **Enable Host Scan/CSD** if it is not already checked.

Step 10 Click **Apply**.

Note If AnyConnect Essentials is enabled on the ASA, you receive a message that CSD will not work with it. You have the choice to **Disable** or **Keep** AnyConnect Essentials.

Step 11 From the File menu, select **Save Running Configuration To Flash**.

Enabling or Disabling Host Scan

When you first install or upgrade a Host Scan image using ASDM, you enable the image as part of that procedure. See "Installing and Enabling Host Scan on the ASA" section on page 87-115.

Otherwise, to enable or disable a Host Scan image using ASDM, follow this procedure:

Step 1 Open ASDM and choose **Configuration > Remote Access VPN > Host Scan Image.** ASDM opens the Host Scan Image panel (Figure 87-7).

Step 2 Check **Enable Host Scan/CSD** to enable Host Scan or uncheck **Enable Host Scan/CSD to disable Host Scan.**

Step 3 Click **Apply**.

Enabling or Disabling CSD on the ASA

Enabling CSD loads the CSD configuration file, data.xml, from the flash device to the running configuration.

Disabling CSD does not alter the CSD configuration.

Use ASDM to enable or disable CSD as follows:

Step 1 Choose **Configuration > Clientless SSL VPN > Secure Desktop > Setup**.

ASDM opens the Setup pane (Figure 87-7).

Note The Secure Desktop Image field displays the image (and version) that is currently installed. The Enable Secure Desktop check box indicates whether CSD is enabled.

Step 2 Check or uncheck **Enable Secure Desktop** and click **Apply**.

ASDM enables or disables CSD.

Step 3 Click the **X** in the upper right corner of the ASDM window to exit.

A window displays the following message:

```
The configuration has been modified. Do you want to save the running configuration to
flash memory?
```

Step 4 Click **Save**. ASDM saves the configuration and closes.

Viewing the Host Scan Version Enabled on the ASA

Open ASDM and select **Configuration > Remote Access VPN > Host Scan Image**.

If there is a Host Scan image designated in the Host Scan Image location field, and the Enable HostScan/CSD box is checked, the version of that image is the Host Scan version being used by the ASA.

If the Host Scan Image filed is empty, and the Enable HostScan/CSD box is checked, select **Configuration > Remote Access VPN > Secure Desktop Manager**. The version of CSD in the Secure Desktop Image Location field is the Host Scan version being used by the ASA.

Uninstalling Host Scan

Uninstalling Host Scan package removes it from view on the ASDM interface and prevents the ASA from deploying it even if Host Scan or CSD is enabled. Uninstalling Host Scan does not delete the Host Scan package from the flash drive.

Uninstall Host Scan on the security appliance as follows:

Step 1 Open ASDM and select **Configuration > Remote Access VPN > Host Scan Image**.

Step 2 In the Host Scan Image pane, click **Uninstall**. ASDM removes the text from the Location text box.

Step 3 From the File menu select **Save Running Configuration to Flash**.

Uninstalling CSD from the ASA

Uninstalling CSD removes the CSD configuration file, data.xml, from the desktop directory on the flash card. If you want to retain the file, copy it using an alternative name or download it to your workstation before you uninstall CSD.

Uninstall CSD on the security appliance as follows:

Step 1 Open ASDM and choose **Configuration > Remote Access VPN > Secure Desktop Manager > Setup**.

ASDM opens the Setup pane (Figure 87-7).

Step 2 Click **Uninstall**.

A confirmation window displays the following message:

```
Do you want to delete disk0:/csd_<n>.<n>.*.pkg and all CSD data files?
```

Step 3 Click **Yes**.

ASDM removes the text from the Location text box and removes the Secure Desktop Manager menu options below Setup.

Step 4 Click the **X** in the upper right corner of the ASDM window to exit.

A window displays the following message:

```
The configuration has been modified. Do you want to save the running configuration to
flash memory?
```

Step 5 Click **Save**. ASDM saves the configuration and closes.

Assigning AnyConnect Posture Module to a Group Policy

Step 1 Open ASDM and choose **Configuration > Remote Access VPN > Network (Client) Access > Group Policies.**

Step 2 In the Group Policies panel, click **Add** to create a new group policy or select the group policy to which you want to assign the Host Scan package and click **Edit**.

Step 3 In the Edit Internal Group Policy panel, expand the **Advanced** navigation tree on the left side of the panel and select **AnyConnect Client**.

Step 4 Uncheck the Optional Client Modules to Download **Inherit** checkbox.

Step 5 In the Optional Client Modules to Download drop down menu, check the AnyConnect Posture Module and click **OK**.

Step 6 Click **OK**.

Other Important Documentation Addressing Host Scan

Once Host Scan gathers the posture credentials from the endpoint computer, you will need to understand subjects like, configuring prelogin policies, configuring dynamic access policies, and using Lua expressions to make use of the information.

These topics are covered in detail in these documents:

- Cisco Secure Desktop Configuration Guides
- Cisco Adaptive Security Device Manager Configuration Guides

See also the Cisco *AnyConnect Secure Mobility Client Administrator Guide, Release 3.0* for more information about how Host Scan works with AnyConnect clients.

Configuring Maximum VPN Sessions

To specify the maximum allowed number of VPN sessions or AnyConnect client VPN sessions, perform the following steps:

Step 1 Choose **Configuration > Remote Access VPN > Advanced > Maximum VPN Sessions**.

Step 2 In the Maximum AnyConnect Sessions field, enter the maximum number of sessions allowed.

Valid values range from 1 to the maximum number of sessions that are allowed by your license.

Step 3 In the Maximum Other VPN Sessions field, enter the maximum number of VPN sessions allowed, which includes Cisco VPN client (IPsec IKEv1) LAN-to-LAN VPN, and clientless SSL VPN sessions.

Valid values range from 1 to the maximum number of sessions that are allowed by your license.

Step 4 Click **Apply**.

Configuring the Pool of Cryptographic Cores

You can change the allocation of cryptographic cores on Symmetric Multi-Processing (SMP) platforms to give you better throughput performance for AnyConnect TLS/DTLS traffic. These changes can accelerate the SSL VPN datapath and provide customer-visible performance gains in AnyConnect, smart tunnels, and port forwarding. To configure the pool of cryptographic cores, perform the following steps.

Limitations

- Cryptographic core rebalancing is available on the following platforms:
  - 5585-X
  - 5580
  - 5545-X

  - 5555-X
  - ASASM
- The large modulus operation is only available for 5510, 5520, 5540, and 5550 platforms.

Detailed Steps

Step 1 Choose **Configuration > Remote Access VPN > Advanced > Crypto Engine**.

Step 2 Check **Enable Large Modulus Acceleration** if you want to perform large modulus operations in the hardware.

Step 3 From the Accelerator Bias drop-down menu, choose one of the following:

Note This field only shows up if the feature is available in ASA.

- **balanced**—Equally distributes cryptography hardware resources (Admin/SSL and IPsec cores).
- **ipsec**—Allocates cryptography hardware resources to favor IPsec (includes SRTP encrypted voice traffic).
- **ssl**—Allocates cryptography hardware resources to favor Admin/SSL.

Step 4 Click **Apply**.

Configuring IP Addresses for VPNs

This chapter describes IP address assignment methods.

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. Once that connection is made, the second set connects client and server through the VPN tunnel.

In ASA 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 VPN 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.

This chapter includes the following sections:

- Configuring an IP Address Assignment Policy, page 88-1
- Configuring Local IP Address Pools, page 88-3
- Configuring AAA Addressing, page 88-5
- Configuring DHCP Addressing, page 88-6

Configuring an IP Address Assignment Policy

The ASA can use one or more of the following methods for assigning IP addresses to remote access clients. If you configure more than one address assignment method, the ASA searches each of the options until it finds an IP address. By default, all methods are enabled.

- **Use authentication server** — Retrieves addresses from an external authentication, authorization, and accounting server on a per-user basis. If you are using an authentication server that has IP addresses configured, we recommend using this method. You can configure AAA servers in the Configuration > AAA Setup pane.This method is available for IPv4 and IPv6 assignment policies.
- **Use DHCP** — Obtains IP addresses from a DHCP server. If you want to use DHCP, you must configure a DHCP server. You must also define the range of IP addresses that the DHCP server can use. If you use DHCP, configure the server in the Configuration > Remote Access VPN > DHCP Server pane. This method is available for IPv4 assignment policies.

- **Use an internal address 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. This method is available for IPv4 and IPv6 assignment policies.
  - Allow the reuse of an IP address so many 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 ASA 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.This configurable element is available for IPv4 assignment policies.

Use one of these methods to specify a way to assign IP addresses to remote access clients.

- Configuring IPv4 and IPv6 Address Assignments using ASDM

Configuring IPv4 and IPv6 Address Assignments using ASDM

Step 1 Select **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Assignment Policy**

Step 2 In the IPv4 Policy area, check the address assignment method to enabled it or uncheck the address assignment method to disable it. These methods are enabled by default:

- Use Authentication server. Enables the use of a Authentication Authorization and Accounting (AAA) server you have configured to provide IP addresses.
- Use DHCP. Enables the use of a Dynamic Host Configuration Protocol (DHCP) server you have configured to provide IP addresses.
- Use internal address pools: Enables the use of a local address pool configured on the ASA.

If you enable **Use internal address pools**, you can also enable the reuse of an IPv4 address after it has been released. You can specify a range of minutes from 0-480 after which the IP v4 address can be reused.

Step 3 In the IPv6 Policy area, check the address assignment method to enable it or uncheck the address assignment method to disable it. These methods are enabled by default:

- Use Authentication server. Enables the use of a Authentication Authorization and Accounting (AAA) server you have configured to provide IP addresses.
- Use internal address pools: Enables the use of a local address pool configured on the ASA.

Step 4 Click **Apply**.

Step 5 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
•	—	•	—	—

Viewing Address Assignment Methods

Use one of these methods to view the address assignment method configured on the ASA:

Viewing IPv4 and IPv6 Address Assignments using ASDM

Select **Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Assignment Policy**

Configuring Local IP Address Pools

To configure IPv4 or IPv6 address pools for VPN remote access tunnels, open ASDM and select **Configuration > Remote Access VPN > Network (Client) Access > Address Management > Address Pools > Add/Edit IP Pool**. To delete an address pool, open ASDM and select **Configuration > Remote Access VPN > Network (Client) Access > Address Management > Address Pools**. Select the address pool you want to delete and click **Delete**.

The ASA uses address pools based on the connection profile or group policy for the connection. The order in which you specify the pools is important. If you configure more than one address pool for a connection profile or group policy, the ASA uses them in the order in which you added them to the ASA.

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.

Use one of these methods to configure a local IP address pool:

- Configuring Local IPv4 Address Pools Using ASDM, page 88-3
- Configuring Local IPv6 Address Pools Using ASDM, page 88-4

Configuring Local IPv4 Address Pools Using ASDM

The IP Pool area shows each configured address pool by name 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 ASA 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.

Step 1 **Select Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools.**

Step 2 To add an IPv4 address, click **Add > IPv4 Address pool**. To edit an existing address pool, select the address pool in the address pool table and click **Edit**.

Step 3 In the Add/Edit IP Pool dialog box enter this information:

- Pool Name—Enter the name of the address pool. It can be up to 64 characters
- Starting Address—Enter the first IP address available in each configured pool. Use dotted decimal notation, for example: 10.10.147.100.
- Ending Address—Enter the last IP address available in each configured pool. User dotted decimal notation, for example: 10.10.147.177.
- Subnet Mask—Identifies the subnet on which this IP address pool resides.

Step 4 Click **Apply**.

Step 5 Click **OK**.

Configuring Local IPv6 Address Pools Using ASDM

The IP Pool area shows each configured address pool by name with a starting IP address range, the address prefix, and the number of addresses configurable in the pool. If no pools exist, the area is empty. The ASA 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.

Step 1 **Select Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools.**

Step 2 To add an IPv6 address, click **Add > IPv6 Address pool**. To edit an existing address pool, select the address pool in the address pool table and click **Edit**.

Step 3 In the Add/Edit IP Pool dialog box enter this information:

- Name—Displays the name of each configured address pool.
- Starting IP Address—Enter the first IP address available in the configured pool. For example: 2001:DB8::1.
- Prefix Length— Enter the IP address prefix length in bits. For example 32 represents /32 in CIDR notation. The prefix length defines the subnet on which the pool of IP addresses resides.
- Number of Addresses—Identifies the number of IPv6 addresses, starting at the Starting IP Address, there are in the pool.

Step 4 Click **Apply**.

Step 5 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
•	—	•	—	—

Configuring AAA Addressing

To use a AAA server to assign addresses for VPN remote access clients, you must first configure a AAA server or server group. See the **aaa-server protocol** command in the command reference and the "Configuring AAA Server Groups" section on page 46-11.

In addition, the user must match a connection profile configured for RADIUS authentication.

The following examples illustrate how to define a AAA server group called RAD2 for the tunnel group named firstgroup. It includes one more step than is necessary, in that previously you might have named the tunnel group and defined the tunnel group type. This step appears in the following example as a reminder that you have no access to subsequent tunnel-group commands until you set these values.

An overview of the configuration that these examples create follows:

```
hostname(config)# vpn-addr-assign aaa
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config-general)# authentication-server-group RAD2
```

To configure AAA for IP addressing, perform the following steps:

Step 1 To configure AAA as the address assignment method, enter the **vpn-addr-assign** command with the **aaa** argument:

```
hostname(config)# vpn-addr-assign aaa
hostname(config)#
```

Step 2 To establish the tunnel group called firstgroup as a remote access or LAN-to-LAN tunnel group, enter the **tunnel-group** command with the **type** keyword. The following example configures a remote access tunnel group.

```
hostname(config)# tunnel-group firstgroup type ipsec-ra
hostname(config)#
```

Step 3 To enter general-attributes configuration mode, which lets you define a AAA server group for the tunnel group called firstgroup, enter the **tunnel-group** command with the **general-attributes** argument.

```
hostname(config)# tunnel-group firstgroup general-attributes
hostname(config-general)#
```

Step 4 To specify the AAA server group to use for authentication, enter the **authentication-server-group** command.

```
hostname(config-general)# authentication-server-group RAD2
hostname(config-general)#
```

This command has more arguments that this example includes. For more information, see the command reference.

Configuring DHCP Addressing

To use DHCP to assign addresses for VPN clients, you must first configure a DHCP server and the range of IP addresses that the DHCP server can use. Then you define the DHCP server on a connection profile basis. Optionally, you can also define a DHCP network scope in the group policy associated with a connection profile or username. This is either an IP network number or IP Address that identifies to the DHCP server which pool of IP addresses to use.

The following examples define the DHCP server at IP address 172.33.44.19 for the connection profile named **firstgroup**. They also define a DHCP network scope of 192.86.0.0 for the group policy called **remotegroup**. (The group policy called remotegroup is associated with the connection profile called firstgroup). If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

The following configuration includes more steps than are necessary, in that previously you might have named and defined the connection profile type as remote access, and named and identified the group policy as internal or external. These steps appear in the following examples as a reminder that you have no access to subsequent tunnel-group and group-policy commands until you set these values.

Guidelines and Limitations

You can only use an IPv4 address to identify a DHCP server to assign client addresses.

Configuring DHCP Addressing Using ASDM

Detailed Steps

Step 1 Select **Configuration > Remote Access VPN > Network (Client) Access> AnyConnect Connection Profiles**.

Step 2 In the Connection Profiles Area click **Add** or **Edit**.

Step 3 Click **Basic** in the configuration tree for the connection profile.

Step 4 In the Client Address Assignment area, enter the IPv4 address of the DHCP server you want to use to assign IP addresses to clients. For example, **172.33.44.19**.

Step 5 Edit the group-policy associated with the connection profile to define the DHCP scope. Select **Configuration > Remote Access VPN > Network (Client) Access> Group Policies.**

Step 6 Double-click the group policy you want to edit.

Step 7 Click **Servers** in the configuration tree.

Step 8 Expand the **More Options** area by clicking the down arrow.

Step 9 Uncheck DHCP Scope **Inherit**.

Step 10 Enter the IP network number or IP Address that identifies to the DHCP server which pool of IP addresses to use. For example, **192.86.0.0**.

Step 11 Click **OK**.

Step 12 Click **Apply**.

Configuring Dynamic Access Policies

This chapter describes how to configure dynamic access policies. It includes the following sections.

- Information About Dynamic Access Policies, page 89-1
- Licensing Requirements for Dynamic Access Policies, page 89-3
- Dynamic Access Policies Interface, page 89-4
- Configuring Dynamic Access Policies, page 89-6
- Testing Dynamic Access Policies, page 89-8
- DAP and Authentication, Authorization, and Accounting Services, page 89-9
- Configuring Endpoint Attributes Used in DAPs, page 89-13
- Configuring DAP Access and Authorization Policy Attributes, page 89-27
- Guide to Creating DAP Logical Expressions using LUA, page 89-31

Information About Dynamic 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 ASA 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 ASA grants access to a particular user for a particular session based on the policies you define. The ASA 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 ASA uses for selecting and applying DAP records during session establishment. Stored on the ASA. You can use ASDM to modify it and upload it to the ASA 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.

Refer to the *Dynamic Access Deployment Guide* (https://supportforums.cisco.com/docs/DOC-1369) for additional information.

DAP and Endpoint Security

The ASA obtains endpoint security attributes by using posture assessment tools that you configure. These posture assessment tools include the AnyConnect posture module, the independent Host Scan package, Cisco Secure Desktop, and NAC.

Table 89-1 identifies each of the remote access protocols DAP supports, the posture assessment tools available for that method, and the information that tool provides.

Table 89-1 DAP Posture Assessment

Remote Access Protocol	AnyConnect Posture Module Host Scan package Cisco Secure Desktop (without Endpoint Assessment Host Scan Extension enabled)	AnyConnect Posture Module Host Scan package Cisco Secure Desktop (with Endpoint Assessment Host Scan Extension enabled)	NAC	Cisco NAC Appliance
	Returns file 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	No	No	Yes	Yes
Cisco AnyConnect VPN	Yes	Yes	Yes	Yes
Clientless VPN	Yes	Yes	No	No
PIX Cut-through Proxy	No	No	No	No

DAP Support for Remote Access Connection Types

The DAP system supports the following remote access methods:

- IPsec VPN
- Clientless (browser-based) SSL VPN
- Cisco AnyConnect Secure Mobility Client (SSL VPN)
- PIX cut-through proxy (posture assessment not available)

Remote Access Connection Sequence with DAPs

The following sequence outlines a typical remote access connection establishment.

1. A remote client attempts a VPN connection.
2. The ASA performs posture assessment, using configured NAC and Cisco Secure Desktop Host Scan values.

3. The ASA authenticates the user via AAA. The AAA server also returns authorization attributes for the user.
4. The ASA applies AAA authorization attributes to the session, and establishes the VPN tunnel.
5. The ASA selects DAP records based on the user AAA authorization information and the session posture assessment information.
6. The ASA aggregates DAP attributes from the selected DAP records, and they become the DAP policy.
7. The ASA applies the DAP policy to the session.

Licensing Requirements for Dynamic Access Policies

The following tables shows the licensing requirements for enforcing Dynamic Access Policies.

Advanced Endpoint Assessment license

Note This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	Advanced Endpoint Assessment License.

SSL VPN license (client)

Note This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	AnyConnect Premium License

AnyConnect Mobile License

Note This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	AnyConnect Mobile license.[1]

1. This license requires activation of one of the following licenses to specify the total number of SSL VPN sessions permitted: AnyConnect Essentials or AnyConnect Premium.

Dynamic Access Policies Interface

Figure 89-1 shows the Dynamic Access Policies pane.

Figure 89-1 Dynamic Access Policies ASDM pane

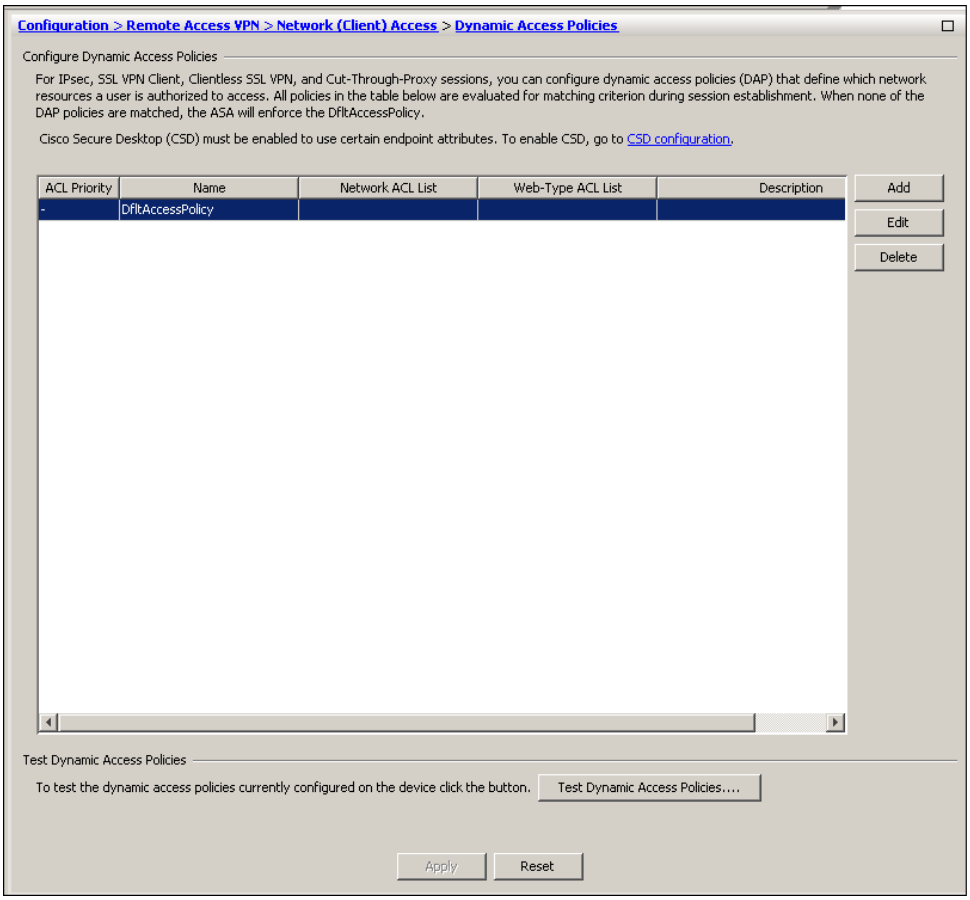

Fields

- ACL Priority—Displays the priority of the DAP record. The ASA uses this value to logically sequence the access lists when aggregating the network and web-type ACLs from multiple DAP records. The ASA 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.
- Find — You can search for a Dynamic Access Policy (DAP) by using the **Find** field. Start typing in the field and the tool will search the beginning characters of every field of the DAP table for a match. You can use wild cards to expand your search.

  For example typing `sal` in the **Find** field will match a DAP named `Sales` but not a DAP named `Wholesalers`. If you type `*sal` in the **Find** field, the search will find the first instance of either Sales or Wholesalers in the table.
- Find Arrows — Use the up and down arrows to skip up or down to the next string match.
- Match Case — Checking the Match Case check box will make your search case-sensitive.

Figure 89-2 shows the Add Dynamic Access Policy pane.

Figure 89-2 Add/Edit Dynamic Access Policies Pane

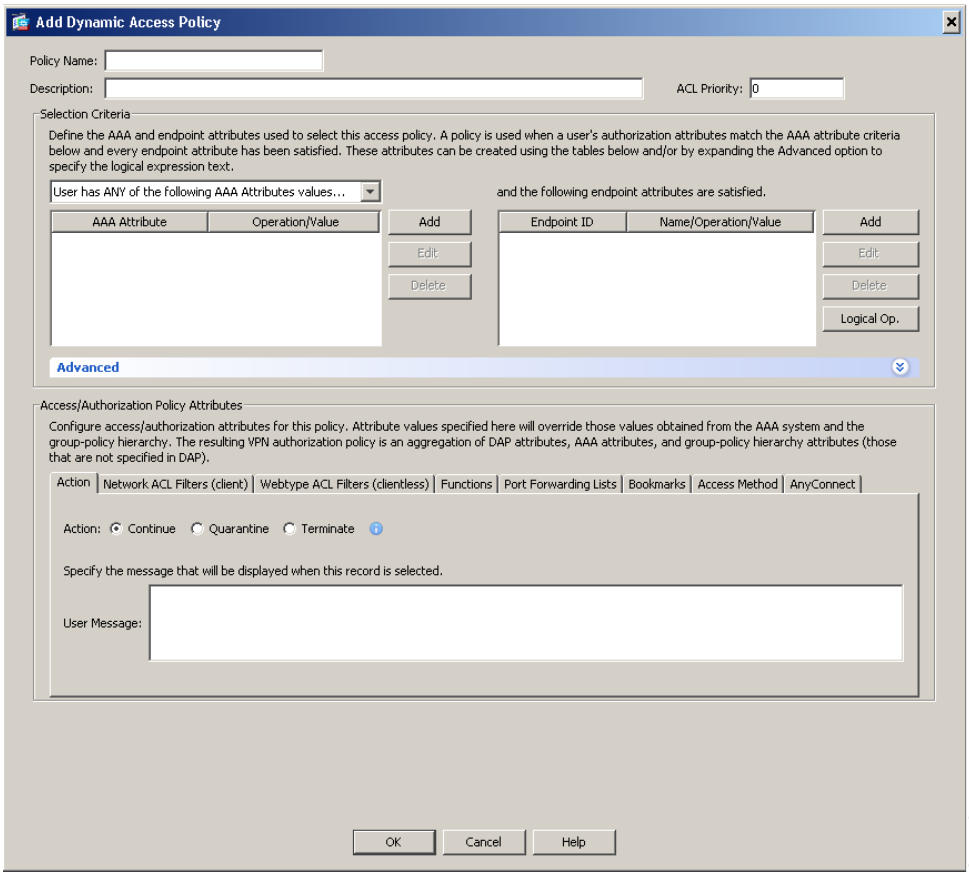

Configuring Dynamic Access Policies

Prerequisites

- Other than where noted, you must install Cisco Secure Desktop or Host Scan before configuring DAP endpoint attributes.
- Before configuring File, Process, and Registry endpoint attributes, configure File, Process, and Registry Basic Host Scan attributes. For instructions, start ASDM and select **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan** and click **Help**.

Guidelines and Limitations

DAP supports only ASCII characters.

Mobile Device Guidelines

ASA administrators will use AnyConnect Mobile Posture DAP Attributes differently depending on the AnyConnect license they have installed. See Adding Mobile Posture Attributes to a DAP, page 89-16 for more information.

Detailed Steps

Step 1 Start ASDM and select **Configuration > Remote Access VPN > Network (Client) Access** or **Clientless SSL VPN Access > Dynamic Access Policies.**

The Configure Dynamic Access Policies pane opens.

Step 2 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 3 To create a new dynamic access policy, click **Add**. To modify an existing policy, click **Edit**.

The Add/Edit Dynamic Access Policy pane opens.

Step 4 At the top of the Add/Edit Dynamic Access Policy pane, provide a name (required) and a description (optional) of this dynamic access policy.

- The **Policy Name** is a string of 4 through 32 characters, no spaces allowed.
- You are allowed a maximum of 80 charaters in the DAP **Description** field.

Step 5 In the **ACL 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. Values of 0 to 2147483647 are valid. The default value is 0.

Step 6 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, 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 ASA always selects it since all selection criteria are satisfied.

Step 7 To Set AAA attributes, click **Add** or **Edit** in the AAA Attributes field. Use one or more of these procedures: See the "Configuring AAA Attributes in a DAP" section on page 89-9 for more information.

Step 8 Use one or more of these procedures to **add** or **edit** endpoint attributes to the DAP policy:

- Adding an Anti-Spyware or Anti-Virus Endpoint Attribute to a DAP, page 89-14
- Adding an Application Attribute to a DAP, page 89-15
- Adding Mobile Posture Attributes to a DAP, page 89-16
- Adding a File Endpoint Attribute to a DAP, page 89-17
- Adding a Device Endpoint Attribute to a DAP, page 89-18
- Adding a NAC Endpoint Attribute to a DAP, page 89-19
- Adding an Operating System Endpoint Attribute to a DAP, page 89-20
- Adding a Personal Firewall Endpoint Attribute to a DAP, page 89-20
- Adding a Policy Endpoint Attribute to a DAP, page 89-21
- Adding a Process Endpoint Attribute to a DAP, page 89-22
- Adding a Registry Endpoint Attribute to a DAP, page 89-23

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 9 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. This feature that requires knowledge of the Lua programming language.

- 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 ASA processes it, discarding any expressions it cannot parse.
- Guide—Click to display online help for creating these logical operations or see Guide to Creating DAP Logical Expressions using LUA, page 89-31.

Step 10 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. 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. See Configuring DAP Access and Authorization Policy Attributes, page 89-27 for more information.

Step 11 Click **OK**.

Tip If you want to test your Dynamic Access Policy, in the Configure Dynamic Access Policies dialog box, click **Test Dynamic Access Policies** and add the attributes to the test interface. See Testing Dynamic Access Policies, page 89-8.

Testing Dynamic Access Policies

Figure 89-3 Test Dynamic Access Policies Pane

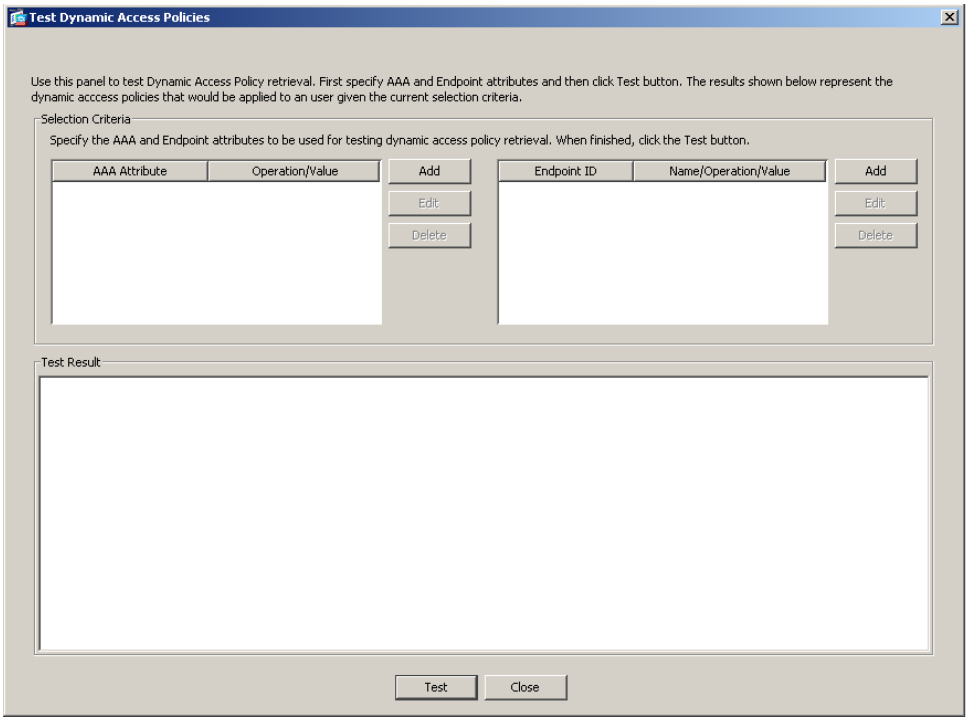

This pane lets you test the retrieval of the set of DAP records configured on the device by specifying authorization attribute value pairs. To specify these pairs, use the Add/Edit buttons associated with the AAA Attribute and Endpoint Attribute tables. The dialogs that display when you click these Add/Edit 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.

DAP and Authentication, Authorization, and Accounting Services

DAP complements AAA services. It provides a limited set of authorization attributes that can override those AAA provides. The ASA selects DAP records based on the AAA authorization information for the user and posture assessment information for the session. The ASA 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 ASA receives from a RADIUS or LDAP server. For more information about DAP and AAA, see the section, Configuring AAA Attributes in a DAP.

Configuring AAA Attributes in a DAP

Figure 89-4 shows the Add AAA Attribute dialog box.

Figure 89-4 Add AAA Attribute Dialog Box

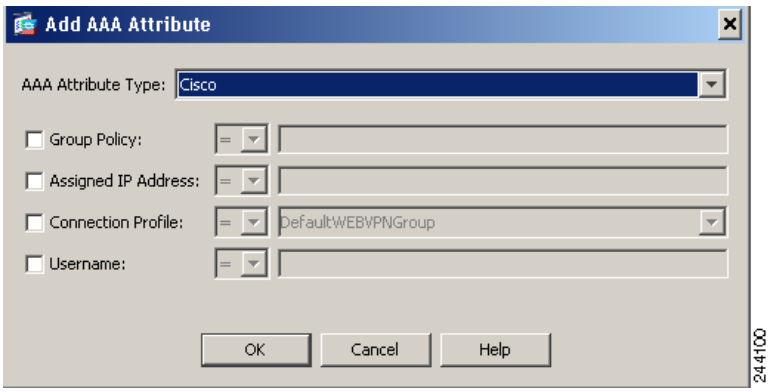

To configure AAA attributes as selection criteria for DAP records, in the Add/Edit AAA Attributes dialog box, set the Cisco, LDAP, or RADIUS attributes that you want to use. You can set these attributes either to = or != the value you enter. There is no limit for the number of AAA attributes for each 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.

  - Assigned IP Address—Enter the IPv4 address you want to specify for the policy. 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.
  - Assigned IPv6 Address—Enter the IPv6 address you want to specify for the policy.
  - 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 ASA 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.

Retrieving Active Directory Groups

Figure 89-5 shows the Retrieve AD Groups from Selected AD Server Group pane.

Figure 89-5 Retrieve AD Groups Dialog Box

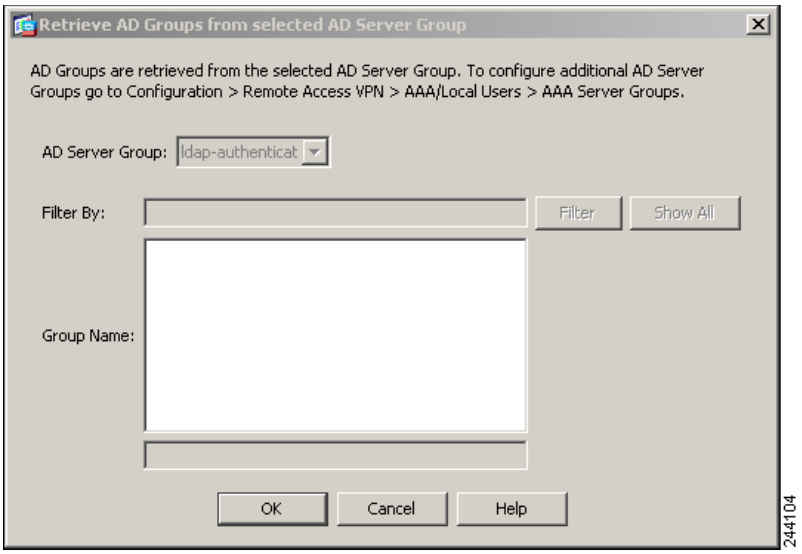

You can query an Active Directory server for available AD groups in this pane. This feature applies only to Active Directory servers using LDAP. Use the group information to specify dynamic access policy AAA selection criteria.

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 ASA 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.

AAA Attribute Definitions

Table 74-2 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.

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Cisco	aaa.cisco.grouppolicy	AAA	string	64	Group policy name on the ASA 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
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.

Configuring Endpoint Attributes Used in DAPs

Endpoint attributes contain information about the endpoint system environment, posture assessment results, and applications. The ASA 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 ASA to select it. The ASA selects only DAP records that satisfy every condition configured.

For detailed information about Endpoint attributes, see Endpoint Attribute Definitions.

Configuring endpoint attributes as selection criteria for DAP records is part of the larger process of Configuring Dynamic Access Policies. Read Configuring Dynamic Access Policies, page 89-6 before you configuring endpoint attributes in DAPs.

This section includes the following topics:

- Adding an Anti-Spyware or Anti-Virus Endpoint Attribute to a DAP, page 89-14
- Adding an Application Attribute to a DAP, page 89-15
- Adding Mobile Posture Attributes to a DAP, page 89-16
- Adding a File Endpoint Attribute to a DAP, page 89-17
- Adding a Device Endpoint Attribute to a DAP, page 89-18
- Adding a NAC Endpoint Attribute to a DAP, page 89-19
- Adding an Operating System Endpoint Attribute to a DAP, page 89-20
- Adding a Personal Firewall Endpoint Attribute to a DAP, page 89-20
- Adding a Policy Endpoint Attribute to a DAP, page 89-21
- Adding a Process Endpoint Attribute to a DAP, page 89-22
- Adding a Registry Endpoint Attribute to a DAP, page 89-23

Figure 89-6 shows the Add Endpoint Attributes dialog box.

Figure 89-6 Add Endpoint Attributes Dialog Box

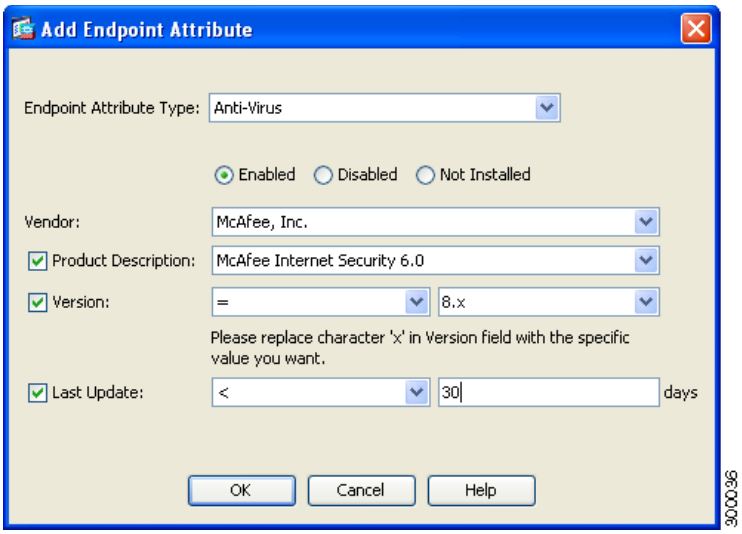

Adding an Anti-Spyware or Anti-Virus Endpoint Attribute to a DAP

Prerequisites

Configuring Anti-Spyware and Anti-Virus endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Anti-Spyware and Anti-Virus endpoint attributes.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match Any** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select Anti-Spyware or Anti-Virus.

Step 2 Click the appropriate **Enabled, Disabled, or Not Installed** button to indicate whether the selected endpoint attribute and its accompanying qualifiers (fields below the Enabled/Disabled/Not Installed buttons) must be enabled, disabled, or are not installed.

Step 3 From the **Vendor ID** list box, click the name of the anti-spyware or anti-virus vendor you are testing for.

Step 4 Check the **Product Description** check box and select from the list box the vendor's product name you are testing for.

Step 5 Check the **Version** checkbox and set the operation field to equal to (=), not equal (!=), less than (<), greater than (>), less that or equal to (<=), or greater than or equal to (>=) the product version number you select from ther **Version** list box.

If the choice in the version list box has an x, such as 3.x, replace the x with a specific release number, for example, 3.5.

Step 6 Check the Last Update check box. 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.

Step 7 Click **OK**.

Step 8 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

- See Endpoint Attribute Definitions, page 89-24 for additional information on the antispyware and antivirus endpoint attribute requirements.
- See DAP and AntiVirus, AntiSpyware, and Personal Firewall Programs, page 89-24 for information on how Host Scan checks for antivirus, antispyware, and personal firewall programs that are memory-resident.

Adding an Application Attribute to a DAP

Prerequisites

Configuring Application endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Application endpoint attributes.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match Any** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Application**.

Step 2 In the Client Type operation field, select equals (=) or does not eaqual (!=).

Step 3 In the Client type list box, indicate the type of remote access connection you are testing for.

Step 4 Click **OK**.

Step 5 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Application endpoint attribute requirements.

Adding Mobile Posture Attributes to a DAP

Licensing

Mobile posture requires an AnyConnect Mobile license and an AnyConnect Premium license installed on the ASA. Enterprises that install these licenses will be able to enforce DAP policies on supported mobile devices based on DAP attributes and other existing endpoint attributes. This includes allowing or denying remote access from a mobile device.

Prerequisites

Configuring mobile posture attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Anti-Spyware and Anti-Virus endpoint attributes.

Guidelines

- These mobile posture attributes can be included in a dynamic access polciy and enforced without installing Host Scan or Cisco Secure Desktop on the endpoint.
- Some mobile posture attributes are relevant to the AnyConnect client running on mobile devices only, some mobile posture attributes are relevant to both AnyConnect clients running on mobile devices and AnyConnect desktop clients.
- When specifying mobile posture attributes and application attributes in a dynamic access policy, they both should be set to AnyConnect.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **AnyConnect**.

Step 2 Check the **Client Version** check box and set the operation field to be equal to (=), not equal to (!=), less than (<), greater than (>), less than or equal to (<=), or greater than or equal to (>=) the AnyConnect client version number you then specify in the **Client Version** field.

You can use this field to evaluate the client version on mobile devices, such as mobile phones and tablets, or desktop and laptop devices.

Step 3 Check the **Platform** check box and set the operation field to be equal to (=), or not equal to (!=) the operating system you then select from the **Platform** list box.

You can use this field to evaluate the operating system on mobile devices, such as mobile phones and tablets, as well as the operating system on desktop and laptop devices. Selecting Apple iOS or Android platforms activates the additional attribute fields for Device Type and Device Unique ID.

Step 4 Check the **Platform Version** check box and set the operation field to be equal to (=), not equal to (!=), less than (<), greater than (>), less than or equal to (<=), or greater than or equal to (>=) the operating system version number you then specify in the **Platform Version** field.

If you want to create a DAP record that contains this attribute, be sure to also specify a Platform in the previous step.

Step 5 If you selected the Platform checkbox and selected the Apple iOS or Android platform, you can check the **Device Type** checkbox. Set the operation field to be equal to (=) or not equal to (!=) the mobile device you then select in the **Device Type** field.

When you specify Android in the **Platform** field, you will be able to pick from a list of supported Android devices in the **Device Type** field. When you specify Apple iOS in the Platform field you will be able to pick from a list of supported Apple devices in the Device Type field. In both cases, the proper Android or Apple iOS device type information is substituted for the device type you choice from the list box.

If you have a supported device which is not listed in the Device Type field, you can enter the Android or Apple iOS device type information in the Device Type field. The most reliable way to obtain the device type information is to install the AnyConnect client on the endpoint and perform a DAP Trace. In the DAP trace results, look for the value of **endpoint.anyconnect.devicetype**. That is the value that you need to enter in the Device Type field.

Step 6 If you selected the Platform checkbox and selected the Apple iOS or Android platform, you can check the **Device Unique ID** checkbox. Set the operation field to be equal to (=) or not equal to (!=) the mobile device's unique ID you then specify in the **Device Unique ID** field.

The Device Unique ID distinguishes individual devices allowing you to set policies for a particular mobile device. To obtain a device's unique ID you will need the device to connect to the ASA and perform a DAP trace. See Performing a DAP Trace, page 89-31 for more information.

Step 7 Click **OK**.

Step 8 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the AnyConnect endpoint attribute requirements.

Adding a File Endpoint Attribute to a DAP

Prerequisites

- Configuring File endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure File endpoint attributes.
- Before configuring a File endpoint attribute, define the file for which you want to scan in the Host Scan window for Cisco Secure Desktop. In ASDM select **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan**. Click **Help** on that page for more information.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match All** is used by default.

Detailed Steps

You only need to configure one AnyConnect attribute in the Add Endpoint Attribute field except where noted.

Step 1 In the **Endpoint Attribute Type** list box, select **File**.

Step 2 Select the appropriate **Exists** or **Does not exist** radio 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.

Step 3 In the **Endpoint ID** list box, choose from the drop-down list the endpoint ID that equates to the file entry for which you want to scan.

The file information is displayed below the Endpoint ID list box.

Step 4 Check the **Last Update** check box and set the operation field to be less than (<) or greater than (>) a certain number of days old. Enter the number of days old in the **days** field.

Step 5 Check the **Checksum** checkbox and set the operation field to be equal to (=) or not equal to (!=) the checksum value of the file you are testing for.

Step 6 Click **Compute CRC32 Checksum** to determine the checksum value of the file you are testing for.

Step 7 Click OK.

Step 8 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the File endpoint attribute requirements.

Adding a Device Endpoint Attribute to a DAP

Prerequisites

Configuring Device endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Device endpoint attributes.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match Any** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Device**.

Step 2 Check the **Host Name** checkbox and set the operation field to be equal to (=) or not equal to (!=) the host name of the device you are testing for. Use the computer's host name only, not the fully qualified domain name (FQDN).

Step 3 Check the **MAC address** checkbox and set the operation field to be equal to (=) or not equal to (!=) the MAC address of the network interface card you are testing for. Only one MAC address per entry. The address must be in the format xxxx.xxxx.xxxx where x is a hexadecimal character.

Step 4 Check the **BIOS Serial Number** checkbox and set the operation field to be equal to (=) or not equal to (!=) the BIOS serial number value of the device you are testing for. The number format is manufacturer-specific. There is no format requirement.

Step 5 Check the **TCP/UDP Port Number** checkbox and set the operation field to be equal to (=) or not equal to (!=) the TCP or UDP port in listening state that you are testing for.

In the TCP/UDP combo box, select the kind of port you are testing for: TCP (IPv4), UDP(IPv4), TCP (IPv6), or UDP (IPv6). If you are testing for more than one port, make several individual endpoint attribute rules in the DAP and specify one port in each.

Step 6 Check the **Privacy Protection** checkbox and set the operation field to be equal to (=) or not equal to (!=) the component CSD uses to execute the PreLogin Policy.

Step 7 Check the **Version of Secure Desktop (CSD)** checkbox and set the operation field to be equal to (=) or not equal to (!=) the version of the Host Scan image running on the endpoint.

Step 8 Check the **Version of Endpoint Assessment** checkbox and set the operation field to be equal to (=) or not equal to (!=) the version of endpoint assessment (OPSWAT) you are testing for.

Step 9 Click OK.

Step 10 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Device endpoint attribute requirements.

Adding a NAC Endpoint Attribute to a DAP

Prerequisites

Configuring NAC endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure NAC endpoint attributes.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match Any** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **NAC**.

Step 2 Check the **Posture Status** checkbox and set the operation field to be equal to (=) or not equal to (!=) the posture token string received by ACS. Enter the posture token string in the Posture Status text box.

Step 3 Click **OK**.

Step 4 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the NAC endpoint attribute requirements.

Adding an Operating System Endpoint Attribute to a DAP

Prerequisites

Configuring **Operating System** endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Operating System endpoint attributes.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Operating System**.

Step 2 Check the **OS Version** checkbox and set the operation field to be equal to (=) or not equal to (!=) the Windows, Mac, or Linux operating system you set in the **OS Version** list box.

Step 3 Check the **OS Update** checkbox and set the operation field to be equal to (=) or not equal to (!=) the Windows, Mac, or Linux service pack for the operating system you enter in the **OS Update** text box.

Step 4 Click **OK**.

Step 5 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Operating System endpoint attribute requirements.

Adding a Personal Firewall Endpoint Attribute to a DAP

Prerequisites

Configuring **Personal Firewall** endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Personal Firewall endpoint attributes.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Operating System**.

Step 2 Click the appropriate **Enabled, Disabled, or Not Installed** button to indicate whether the selected endpoint attribute and its accompanying qualifiers (fields below the Enabled/Disabled/Not Installed buttons) must be enabled, disabled, or are not installed.

Step 3 From the **Vendor ID** list box, click the name of the personal firewall vendor you are testing for.

Step 4 Check the **Product Description** check box and select from the list box the vendor's product name you are testing for.

Step 5 Check the **Version** checkbox and set the operation field to equal to (=), not equal (!=), less than (<), greater than (>), less that or equal to (<=), or greater than or equal to (>=) the product version number you select from ther **Version** list box.

If the choice in the **Version** list box has an x, such as 3.x, replace the x with a specific release number, for example, 3.5.

Step 6 Click **OK.**

Step 7 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

- See Endpoint Attribute Definitions, page 89-24 for additional information on the Personal Firewall endpoint attribute requirements.
- See DAP and AntiVirus, AntiSpyware, and Personal Firewall Programs, page 89-24 for information on how Host Scan checks for antivirus, antispyware, and personal firewall programs that are memory-resident.

Adding a Policy Endpoint Attribute to a DAP

Prerequisites

Configuring **Policy** endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure **Policy** endpoint attributes.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match Any** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Policy**.

Step 2 Check the **Location** checkbox and set the operation field to be equal to (=) or not equal to (!=) the Cisco Secure Desktop Microsoft Windows location profile. Enter the Cisco Secure Desktop Microsoft Windows location profile string in the **Location** text box.

Step 3 Click **OK**.

Step 4 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Policy endpoint attribute requirements.

Adding a Process Endpoint Attribute to a DAP

Prerequisites

- Configuring **Process** endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Personal Firewall endpoint attributes.
- Before configuring a Process endpoint attribute, define the process for which you want to scan in the Host Scan window for Cisco Secure Desktop. In ASDM select **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan**. Click **Help** on that page for more information.

Guidelines

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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match All** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Process**.

Step 2 Click the appropriate **Exists or Does not exist** button to indicate whether the selected endpoint attribute and its accompanying qualifiers (fields below the Exists and Does not exist buttons) should be present or not.

Step 3 In the **Endpoint ID** list box, choose from the drop-down list the endpoint ID for which you want to scan.

The endpoint ID process information is displayed below the list box.

Step 4 Click **OK**.

Step 5 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Process endpoint attribute requirements.

Adding a Registry Endpoint Attribute to a DAP

Prerequisites

- Configuring **Process** endpoint attributes as selection criteria for DAP records is part of a larger process. Read Configuring Dynamic Access Policies, page 89-6 before you configure Personal Firewall endpoint attributes.
- Before configuring a Registry endpoint attribute, define the registry key for which you want to scan in the Host Scan window for Cisco Secure Desktop. In ASDM select **Configuration > Remote Access VPN > Secure Desktop Manager > Host Scan**. Click **Help** on that page for more information.

Guidelines

- You can only scan for registry endpoint attributes on Windows operating systems.
- 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, after you have defined all instances of the endpoint attribute, click the **Logical Op.** button and select the **Match Any** or **Match All** button. If you do not specify a Logical Operation, **Match All** is used by default.

Detailed Steps

Step 1 In the **Endpoint Attribute Type** list box, select **Registry**.

Step 2 Click the appropriate **Exists or Does not exist** button to indicate whether the **Registry** endpoint attribute and its accompanying qualifiers (fields below the Exists and Does not exist buttons) should be present or not.

Step 3 In the **Endpoint ID** list box, choose from the drop-down list the endpoint ID that equates to the registry entry for which you want to scan.

The registry information is displayed below the Endpoint ID list box.

Step 4 Check the **Value** checkbox and set the operation field to be equal to (=) or not equal to (!=).

Step 5 In the first **Value** list box, identify the registry key as a dword or a string.

Step 6 In the second Value operation list box, enter the value of the resgistry key you are scanning for.

Step 7 If you want to disregard the case of the registry entry when scanning, click the **Caseless** checkbox. If you want the search to be case-sensitive, do not ccheck the Caseless check box.

Step 8 Click **OK**.

Step 9 Return to Configuring Dynamic Access Policies, page 89-6.

Additional References

See Endpoint Attribute Definitions, page 89-24 for additional information on the Registry endpoint attribute requirements.

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.

Endpoint Attribute Definitions

Table 89-2 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 89-2 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 89-2 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
AnyConnect (Does not require Cisco Secure Desktop or Host Scan.)	endpoint.anyconnect.clientversion	Endpoint	version	—	AnyConnect client version.
	endpoint.anyconnect.platform		string	—	Operating system on which AnyConnect client is installed.
	endpoint.anyconnect.platformversion		version	64	Version of operating system on which AnyConnect client is installed.
	endpoint.anyconnect.devicetype		string	64	Mobile device type on which AnyConnect client is installed.
	endpoint.anyconnect.deviceuniqueid		caseless	64	Unique ID of mobile device on which AnyConnect client is installed.
Application	endpoint.application.clienttype	Application	string	—	Client type: CLIENTLESS ANYCONNECT IPSEC L2TP

Table 89-2 Endpoint Attribute Definitions (continued)

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Device	endpoint.device.hostname	Endpoint	string	64	Host Name only. Not FQDN.
	endpoint.device.MAC		string	Must be in the format xxxx.xxxx.xxxx where x is a hexadecimal character.	Mac Address for a network interface card. Only one Mac address per entry.
	endpoint.device.id		string	64	BIOS Serial Number. The number format is manufacturer-specific. There is no format requirement.
	endpoint.device.port		string	An integer between 1 and 65535.	TCP port in listening state. You can define a single port per line.
	endpoint.device.protection		None (Host Scan) Scure Desktop (either Cache Cleaner or Vault)	64	Defines which component of CSD will execute for the particular PreLogin Policy.
	endpoint.device.protection_version		string	64	Version of Host Scan image they are running.
	endpoint.device.protection_extension		string	64	Versionof Endpoint Assessment (OPSWAT)
File	endpoint.file["*label*"].exists	Secure Desktop	true	—	The files exists
	endpoint.file[“label”].endpointid				
	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

Table 89-2 Endpoint Attribute Definitions (continued)

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
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

Configuring DAP Access and Authorization Policy Attributes

To Configure Access and Authorization Policy Attributes for a DAP, click each tab and configure the fields.

- Action Tab—Specifies special processing to apply to a specific connection or session.
  - Continue—(Default) Click to apply access policy attributes to the session.
  - Quarantine—Through the use of quarantine, you can restrict a particular client who already has an established tunnel through a VPN. ASA applies restricted ACLs to a session to form a restricted group, based on the selected DAP record. When an endpoint is not compliant with an administratively defined policy, the user can still access services for remediation (such as updating the antivirus and so on), but restrictions are placed upon the user. After the remediation occurs, the user can reconnect, which invokes a new posture assessment. If this assessment passes, the user connects.

    Note This parameter requires an AnyConnect release that supports AnyConnect Secure Mobility features.

  - 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 490 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 Instructions 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 ASA 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. This field supports unified ACLs which can define access rules for IPv4 and IPv6 network traffic.
  - **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 ASA unless you first delete it from DAP records.
- **Web-Type ACL Filters (clientless) 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 ASA 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 ASA 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 ASA 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 ASA to the destination web server is not secured.

In a clientless VPN connection, the ASA 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 ASA 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 ASA 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 ASA unless you first delete it from DAP records.
- **Bookmarks Tab**—Lets you select and configure bookmarks for certain user session URLs.
  - **Enable bookmarks**—Click to enable. When unchecked, no bookmarks display in the portal page for the connection.
  - **Bookmark** drop-down list—select already configured bookmarks to add to the DAP record.
  - **Manage...**—Click to add, import, export, and delete bookmarks.
  - **Bookmarks (unlabeled)**—Displays the URL lists for the DAP record.
  - **Add>>**—Click to add the selected bookmark from the drop-down list to the URL area on the right.
  - **Delete**—Click to delete the selected bookmark from the URL list area. You cannot delete a bookmark from the ASA 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.
- **AnyConnect Tab**—Lets you choose the status of the Always-on VPN flag.
  - Always-On VPN for AnyConnect client—Determine if the always-on VPN flag setting in the AnyConnect service profile is unchanged, disabled, or if the AnyConnect profile setting should be used.

Note This parameter requires a release of the Cisco IronPort Web Security appliance that provides Secure Mobility Solution licensing support for the Cisco AnyConnect VPN client. It also requires an AnyConnect release that supports "Secure Mobility Solution" features. Refer to the *Cisco AnyConnect VPN Client Administrator Guide* for additional information.

Performing a DAP Trace

By performing a DAP trace you can display the DAP endpoint attributes for all connected devices.

Prerequisites

Log on to the ASA from an SSH terminal and enter Privileged Exec mode. In Privileged Exec mode, the ASA displays this prompt: `hostname#`

Detailed Steps

	Command	Purpose

Step 1	**debug dap trace**
	
	Example
	`hostname# debug dap trace`
		Enables DAP debugs to display all DAP attributes for the session in the terminal window.
		
		Example output:
		This is a small fragment of the output one receives from running the debug dap trace command
		`endpoint.anyconnect.clientversion="0.16.0021";`
		`endpoint.anyconnect.platform="apple-ios";`
		`endpoint.anyconnect.platformversion="4.1";`
		`endpoint.anyconnect.devicetype="iPhone1,2";`
		`endpoint.anyconnect.deviceuniqueid="dd13ce3547f2fa1b2c3d4e5f6g7h8i9j0fa03f75";`

Additional References

In order to search the output of the DAP trace, send the output of the command to a sytem log. To learn more about logging on the ASA see Configuring Logging in the *Cisco ASA 5500 Series Configuration Guide using the CLI, 8.4*.

Guide to Creating DAP Logical Expressions using LUA

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 ASA 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 ASA 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 74-1.
- For a list of endpoint selection attributes, including proper name syntax for creating logical expressions, see Table 89-2.

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 74-1 and Table 89-2 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 ASA 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 ASA 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 ASA 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 ASA 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 89-3 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 89-3 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 ASA 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 ASA. In addition, the ASA 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.

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 ASA for Clientless SSL VPN. For further information see Configuring Clientless SSL VPN Access, page 91-6.

Requiring Usernames and Passwords

Depending on your network, during a remote session users may 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 may have to authenticate in many different contexts, requiring different information, such as a unique username, password, or PIN. Ensure users have the required access.

Table 90-1 lists the type of usernames and passwords that Clientless SSL VPN users might need to know.

Table 90-1 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

Table 90-1 Usernames and Passwords to Give to Clientless SSL VPN Users (continued)

Login Username/ Password Type	Purpose	Entered When
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 ASA 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 ASA to the destination web server is not secure.

Configuring Remote Systems to Use Clientless SSL VPN Features

Table 90-2 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 email via Port Forwarding, Web Access, or Email Proxy

Table 90-2 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 90-2 organizes information by user activity, so you can skip over the information for unavailable features.

Table 90-2 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 8 • Firefox 8 On Linux: • Firefox 8 On Mac OS X: • Safari 5 • Firefox 8
	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 ASA (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 90-2 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 ASA 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 90-2 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 90-2 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 Oracle Java Runtime Environment (JRE) 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.
	Oracle 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 90-2 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 8.x or higher, or Firefox 8.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 ASA 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 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 ASA***/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.

Configuring Clientless SSL VPN

This chapter describes how to configure clientless SSL VPN and includes the following sections:

- Information About Clientless SSL VPN, page 91-2
- Licensing Requirements, page 91-2
- Prerequisites for Clientless SSL VPN, page 91-5
- Guidelines and Limitations, page 91-5
- Configuring Clientless SSL VPN Access, page 91-6
- Configuring the Setup for Cisco Secure Desktop, page 91-13
- Configuring Application Profile Customization Framework, page 91-15
- Using Auto Signon, page 91-25
- Configuring Session Settings, page 91-28
- Java Code Signer, page 91-28
- Encoding, page 91-29
- Content Cache, page 91-30
- Content Rewrite, page 91-31
- Configuring Browser Access to Plug-ins, page 91-33
- Understanding How KCD Works, page 91-38
- Configuring Application Access, page 91-45
- Configuring Port Forwarding, page 91-55
- Application Access User Notes, page 91-64
- Configuring File Access, page 91-68
- Ensuring Clock Accuracy for SharePoint Access, page 91-69
- Customizing the Clientless SSL VPN User Experience, page 91-69
- Using Clientless SSL VPN with PDAs, page 91-74
- Using E-Mail over Clientless SSL VPN, page 91-74
- Configuring Portal Access Rules, page 91-75
- Clientless SSL VPN End User Setup, page 91-77
- Configuring Browser Access to Client-Server Plug-ins, page 91-118
- Customizing the AnyConnect Client, page 91-131

- Configuring Bookmarks, page 91-140
- Sending an Administrator's Alert to Clientless SSL VPN Users, page 91-149

Information About Clientless SSL VPN

Note When the ASA is configured for clientless SSL VPN, you cannot enable security contexts (also called firewall multimode) or Active/Active stateful failover. Therefore, these features become unavailable.

Clientless SSL VPN lets users establish a secure, remote-access VPN tunnel to an ASA using a web browser. Users do not need a software or hardware client.

Clientless SSL VPN provides secure and easy access to a broad range of web resources and both web-enabled and legacy applications from almost any computer that can reach HTTP Internet sites. They include:

- Internal websites
- Web-enabled applications
- NT/Active Directory file shares
- E-mail proxies, including POP3S, IMAP4S, and SMTPS
- Microsoft Outlook Web Access Exchange Server 2000, 2003, and 2007
- Microsoft Web App to Exchange Server 2010 in 8.4(2) and later.
- Application Access (that is, smart tunnel or port forwarding access to other TCP-based applications)

Clientless SSL VPN uses Secure Sockets Layer Protocol and its successor, Transport Layer Security (SSL/TLS1) to provide the secure connection between remote users and specific, supported internal resources that you configure at an internal server. The ASA 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 resources by users of clientless SSL VPN sessions on a group basis. Users have no direct access to resources on the internal network.

Licensing Requirements

The following table shows the licensing requirements for this feature:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1,2]
ASA 5505	AnyConnect Premium license: • Base License or Security Plus license: 2 sessions. • *Optional permanent or time-based licenses: 10 or 25 sessions.* • *Shared licenses are not supported.*[3]
ASA 5510	AnyConnect Premium license: • Base and Security Plus License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, or 250 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5520	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, or 750 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5540	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, or 2500 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5550	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, or 5000 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5580	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, 5000, or 10000 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5512-X	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, or 250 sessions.* • *Optional Shared licenses*[3]*: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*

Model	License Requirement[1,2]
ASA 5515-X	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, or 250 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5525-X	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, or 750 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5545-X	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, or 2500 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5555-X	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, or 5000 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5585-X with SSP-10	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, or 5000 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASA 5585-X with SSP-20, -40, and -60	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, 5000, or 10000 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*
ASASM	AnyConnect Premium license: • Base License: 2 sessions. • *Optional permanent or time-based licenses: 10, 25, 50, 100, 250, 500, 750, 1000, 2500, 5000, or 10000 sessions.* • *Optional Shared licenses[3]: Participant or Server. For the Server license, 500-50,000 in increments of 500 and 50,000-545,000 in increments of 1000.*

1. 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.
2. The maximum combined VPN sessions of *all* types cannot exceed the maximum sessions shown in this table.
3. A shared license lets the ASA act as a shared license server for multiple client ASAs. The shared license pool is large, but the maximum number of sessions used by each individual ASA cannot exceed the maximum number listed for permanent licenses.

Prerequisites for Clientless SSL VPN

See the *Supported VPN Platforms, Cisco ASA 5500 Series* for the platforms and browsers supported by ASA Release 9.0.

Guidelines and Limitations

- ActiveX pages require that you enable ActiveX Relay or enter **activex-relay** 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.
- The ASA does not support clientless access to Windows Shares (CIFS) Web Folders from Windows 7, Vista, Internet Explorer 8-9, Mac OS X, and Linux.
- Certificate authentication, including the DoD Common Access Card and SmartCard, works with the Safari keychain only.
- The ASA does not support DSA or RSA certificates for clientless SSL VPN connections.
- Some domain-based security products have requirements above those requests that originate from the ASA.
- Inspecting configuration control and other inspection features under the Modular Policy Framework are not supported.
- Neither NAT or PAT is applicable to the client.
- Some components of Clientless SSL VPN require the Java Runtime Environment (JRE).With Mac OS X v10.7 and later Java is not installed by default. For details of how to install Java on Mac OS X see http://java.com/en/download/faq/java_mac.xml.
- If you have several group policies configured for the clientless portal, they are displayed in a drop-down on the logon page. If the top of the list of group policies is one that requires a certificate, then as soon as the user gets to the logon page, they must have a matching certificate. If not all your group policies use certificates, then configure the list to display a non-certificate policy first. Name your group polices to sort alphabetically, or prefix them with numbers so an AAA policy shows up first. For example, 1-AAA, 2-Certificate. Or, create a "dummy" group policy named Select-a-Group, and make sure that shows up first.

Observing Clientless SSL VPN Security Precautions

By default, the ASA permits all portal traffic to all web resources (e.g., HTTPS, CIFS, RDP, and plug-ins). The ASA 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, 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.

Figure 91-1 Example URL Typed by User

Figure 91-2 Same URL Rewritten by Security Appliance and displayed on the Browser Window

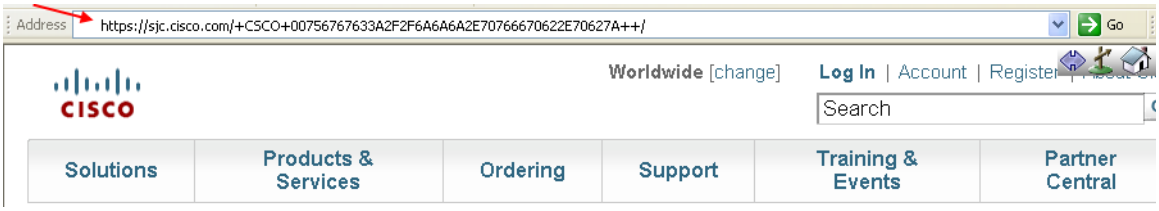

Detailed Steps

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**.

Step 3 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.

Step 4 Assign the web ACL to any policies (group policies, dynamic access policies, or both) that you have configured for clientless access. To assign a web ACL to a DAP, edit the DAP record, and select the web ACL on the **Network ACL Filters** tab.

Step 5 Disable URL entry on the *portal page*, the page that opens upon the establishment of 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. To disable URL entry on a DAP, use ASDM to edit the DAP record, click the **Functions** tab, and check **Disable** next to URL Entry

Step 6 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.

Configuring Clientless SSL VPN Access

When configuring Clientless SSL VPN access, you can do the following:

- Enable or disable ASA interfaces for clientless SSL VPN sessions.
- Choose a port for clientless SSL VPN connections.
- Set a maximum number of simultaneous clientless SSL VPN sessions.

Detailed Steps

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies** pane to configure or create a group policy for clientless access.

Step 2 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles**.

a. Enable or disable **Allow Access** for each ASA interface.

The Interface columns lists the configured interfaces. The WebVPN Enabled field displays the 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.

b. Click the **Port Setting** button, and 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.

Step 3 Navigate to **Configuration > Remote Access VPN > Advanced > Maximum VPN Sessions**, and enter the maximum number of clientless SSL VPN sessions you want to allow in the Maximum Other VPN Sessions field. 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.

Clientless SSL VPN Server Certificate Verification

When connecting to a remote SSL-enabled server through clientless SSL VPN, it is important to know that you can trust the remote server, and that it is in fact the server you are trying to connect to. ASA 9.0 introduces support for SSL server certificate verification against a list of trusted certificate authority (CA) certificates for clientless SSL VPN.

When you connect to a remote server via a web browser using the HTTPS protocol, the server will provide a digital certificate signed by a CA to identify itself. Web browsers ship with a collection of CA certificates which are used to verify the validity of the server certificate. This is a form of public key infrastructure (PKI).

Just as browsers provide certificate management facilities, so does the ASA in the form of trusted certificate pool management facility: trustpools. This can be thought of as a special case of trustpoint representing multiple known CA certificates. The ASA includes a default bundle of certificates, similar to that provided with web browsers, but it is inactive until activated by the administrator.

Note If you are already familiar with trustpools from Cisco IOS then you should be aware that the ASA version is similar, but not identical.

Enabling HTTP Server Verification

To enable HTTPS Server Verification for Clientless SSL VPN users:

Step 1 In the ASDM, navigate to **Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool**.

Figure 91-3 Enabling HTTPS Server Verification in the ASDM

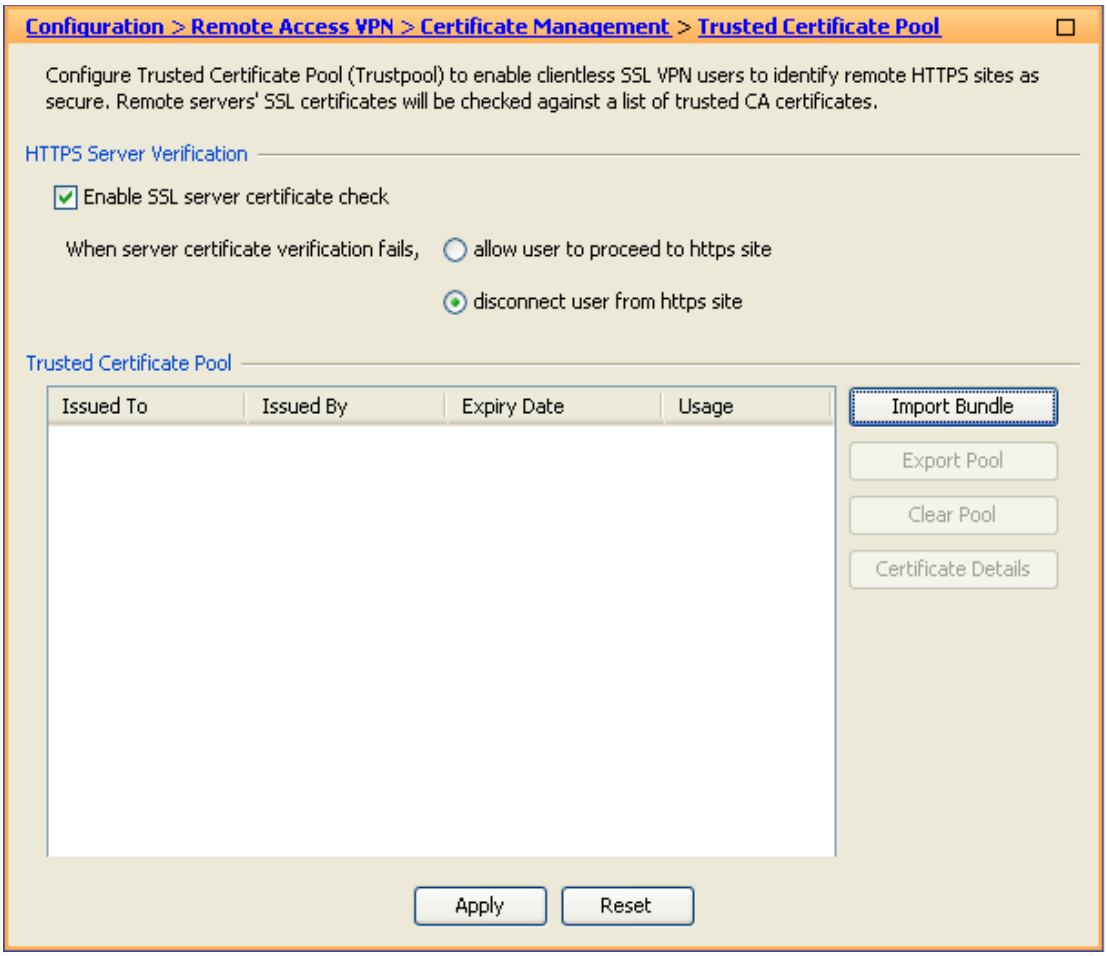

Step 2 Select the **Enable SSL Certificate check** check box.

Step 3 You must decide which action you want to be taken if server certificate verification fails. Click **disconnect user from https site** to disconnect if the server could not be verified. Alternatively, click **allow user to proceed to https site** to allow the user to continue the connection, even if the check failed.

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

Importing a Certificate Bundle

You can import individual certificates or bundles of certificates from a variety of locations in one of the following formats:

- x509 certificates in DER format wrapped in a pkcs7 structure

- a file of concatenated x509 certificates in PEM format (complete with PEM header)

To import a certificate or bundle:

Step 1 In the ASDM, navigate to **Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool**.

Step 2 Click **Import Bundle**.

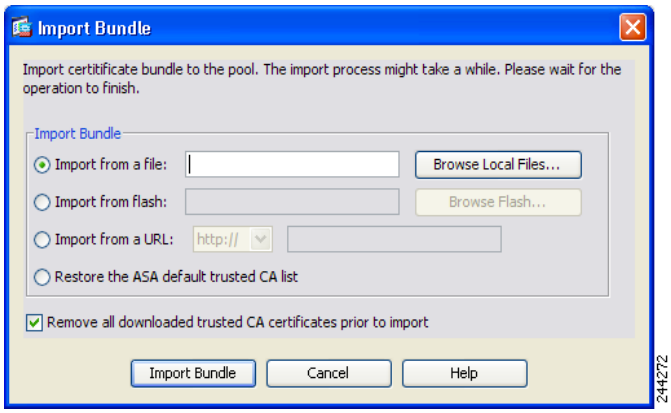

Step 3 Select the location of the bundle:

- If the bundle is stored on your machine, click **Import from a file**, then click **Browse Local Files** and navigate to the bundle.
- If the bundle is stored on the ASA flash file system, click **Import from flash**, then click **Browse Flash** and navigate to the file.
- If the bundle is hosted on a server, click **Import from a URL**, select the protocol from the list, and enter the URL in the box.

Step 4 Click **Import Bundle**. Alternatively, click **Cancel** to abandon your changes.

Note You can select the **Remove all downloaded trusted CA certificates prior to import** check box to clear the trustpool before importing a new bundle.

Exporting the Trustpool

When you have correctly configured the Trustpool you should export the pool. This will enable you to restore the Trustpool to this point, for example if you wish to remove a certificate that was added to the trustpool after the export. You can export the pool to the ASA flash file system or your local file system.

In the ASDM, navigate to **Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool**, then click **Export Pool**.

To export to the local file system:

Step 1 Click **Export to a file**.

Step 2 Click **Browse Local Files.**

Step 3 Navigate to the folder where you want to save the trustpool.

Step 4 Enter a unique memorable name for the trustpool in the **File name** box.

Step 5 Click **Select**.

Step 6 Click **Export Pool** to save the file. Alternatively, click **Cancel** to stop saving.

Removing Certificates

To remove all certificates, in the ASDM, navigate to **Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool**, then click **Clear Pool**.

Note Before clearing the trustpool you should export the current trustpool to enable you to restore your current settings.

Restoring the Default Trusted Certificate Authority List

To restore the default trusted Certificate Authority (CA) list, in the ASDM, navigate to **Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool**, then click **Restore default trusted CA list** and click **Import Bundle**.

Updating the Trustpool

The trustpool should be updated if either of the following conditions exists:

- Any certificate in the trustpool is due to expire or has been re-issued.
- The published CA certificate bundle contains additional certificates that are required by a specific application.

A full update will replace all the certificates in the trustpool.

A practical update enables you to add new certificates or replace existing certificates.

Removing a Certificate Bundle

Clearing the trustpool removal will remove all certificates that are not part of the default bundle.

You cannot remove the default bundle. To clear the trustpool, in the ASDM, navigate to Configuration > Remote Access VPN > Certificate Management > Trusted Certificate Pool, then click Clear Pool.

Note

	Command	Purpose
Step 1	`webvpn`	Switches to group policy webvpn configuration mode.
Step 2	`url-entry disable`	Disables URL entry.

Configuring ACLs

ACLs constrain user access to specific networks, subnets, hosts, and web servers. The Web ACLs table displays the filters configured on the ASA application to the 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 to specific networks, subnets, hosts, and web servers. Each ACE specifies one rule that serves the function of the ACL.

Guidelines

If you do not define any filters, all connections are permitted.

Restrictions

- The ASA 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 ASA denies it. ACEs are referred to as rules in this topic.

Detailed Steps

Web ACLs are configured on the page **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Web ACLs**.

Step 1 Click **Add ACL** to add an ACL or ACE. To insert a new ACE before or after an existing ACE, click **Insert** or **Insert After**.

Step 2 Click **Edit** to highlight the ACE you want to change.

Step 3 Highlight the ACL or ACE you want to remove and click **Delete**. When you delete an ACL, you must delete all of its ACEs. No warning is provided, and or undelete.

Step 4 Use the **Move Up** and **Move Down** buttons to change the order of ACLs or ACEs. The ASA 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.

Step 5 Click **+** to expand or **-** to collapse the list of ACEs under each ACL. The priority of the ACEs under each ACL is displayed. The order in the list determines priority.

Step 6 (Optional) Click **Find** to search for a web ACL. Start typing in the field, and the tool searches the beginning characters of every field for a match. You can use wild cards to expand your search. For example, typing *sal* in the Find field matches a web ACL named sales but not a customization object named wholesalers. If you type **sal* in the Find field, the search finds the first instance of either sales or wholesalers in the table.

Step 7 Use the up and down arrows to skip up or down to the next string match. Check the **Match Case** checkbox to make your search case sensitive.

Step 8 (Optional) Highlight a web ACL and click **Assign** to assign the selected web ACL to one or more VPN group policies, dynamic access policies, or user policies.

Step 9 When you create an ACE, by default it is enabled. Clear the check box to disable an ACE.

The IP address or URL of the application or service to which the ACE applies is displayed. The TCP service to which the ACE applies is also displayed. The Action field displays whether the ACE permits or denies clientless SSL VPN access. The time range associated with the ACE and the logging behavior (either disabled or with a specified level and time interval) is also displayed.

Adding or Editing ACEs

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.

Detailed Steps

Step 1 Permit or deny access to specific networks, subnets, hosts, and web servers specified in the Filter group field.

Step 2 Specify 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.

Step 3 The Rule Flow Diagram graphically depicts the traffic flow using the filter. This area may be hidden.

Step 4 Specify 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 ASA 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.

Configuration Examples for ACLs for Clientless SSL VPN

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.example.com/ directory/file.html	Denies access to the specified file.
Permit	url https://www.example.com/directory	Permits access to the specified location
Deny	url http://*:8080/	Denies HTTPS access to anywhere via port 8080.

Action	Filter	Effect
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.

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 ASA, 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 ASA.

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 ASA, 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 ASA, 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.

Detailed Steps

The Cisco Secure Desktop image loaded into the running configuration is displayed in the Location field. By default, the filename is in the format securedesktop_asa_<n>_<n>*.pkg.

Step 1 Click **Browse Flash** to insert or modify the value in this field.

Step 2 Click **Enable Secure Desktop** 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.

Uploading Images

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 ASA. Use this window to install or upgrade Cisco Secure Desktop.

Prerequisites

- 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.

Detailed Steps

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**.

Step 1 Specify 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.

Step 2 Specify the destination path on the flash device of the ASA 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

The file name of the Cisco Secure Desktop image that you selected on your local computer is displayed in the Browse Flash dialog box. 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.

Configuring Application Profile Customization Framework

Clientless SSL VPN includes an Application Profile Customization Framework option that lets the ASA 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.

Multiple APCF profiles can run in parallel on an ASA. Within an APCF profile script, multiple APCF rules can apply. In this case, the ASA processes the oldest rule first (based on configuration history), then the next oldest rule, and so forth.

You can configure multiple APCF profiles on an ASA. Within an APCF profile script, multiple APCF rules can apply. The ASA processes the oldest rule first, based on configuration history, the next oldest rule next, and so forth.

You can store APCF profiles on the ASA flash memory, or on an HTTP, HTTPS, or TFTP server.

Restrictions

We recommend that you configure an APCF profile only with the assistance of Cisco personnel.

Managing APCF Profiles

You can store APCF profiles on the ASA 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.

Step 1 Navigate to Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Application Helper, where you can perform the following functions.

- Click **Add/Edit** to create a new APCF profile or change an existing one.
  - Select **Flash file** to locate an APCF file stored on the ASA flash memory.

    Then click **Upload** to get an APCF file from a local computer to the ASA flash file system, or Browse to upload select an APCF file that is already in flash memory.
  - Select URL to retrieve the APCF file from an HTTP, HTTPS, FTP, or TFTP server.
- Click **Delete** to remove an existing APCF profile. No confirmation or undo exists.
- Click **Move Up** or **Move Down** to rearrange APCF profiles within the list. The order determines which the APCF profile is used.

Step 2 Click **Refresh** if you do not see the changes you made in the list.

Uploading APCF Packages

Detailed Steps

Step 1 The path to the APCF file on your computer is shown. Click **Browse Local** to automatically insert the path in this field, or enter the path.

Step 2 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.

Step 3 The path on the ASA to upload the APCF file is shown in the Flash File System Path. Click **Browse Flash** to identify the location on the ASA to which you want to upload the APCF file. The Browse Flash dialog box displays the contents of flash memory.

Step 4 The file name of the APCF file you selected on your local computer is displayed. 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.

Step 5 Click **Upload File** when you have identified the location of the APCF file on your computer, and the location where you want to download it to the ASA.

Step 6 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**.

Step 7 Close the Upload Image dialog window. Click **Close** 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**.

Managing APCF Packets

Step 1 Use the following commands to add, edit, and delete APCF packets and put them in priority order:

- APCF File Location—Displays information about the location of the APCF package. This can be on the ASA 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 ASA attempts to use APCF profiles.

Step 2 Click **Flash file** to locate an APCF file stored on the ASA flash memory.

Step 3 Enter the path to an APCF file stored on flash memory. If you already added a path, it displays to an APCF file stored on flash memory after you browse to locate it.

Step 4 Click Browse Flash 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 ASA 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.

APCF Syntax

APCF profiles use XML format, and sed script syntax, with the XML tags in Table 91-1.

Guidelines

Misuse of an APCF profile can result in reduced performance and undesired rendering of content. In most cases, Cisco Engineering supplies APCF profiles to solve specific application rendering issues.

Table 91-1 APCF XML Tags

Tag	Use
<APCF>...</APCF>	The mandatory root element that opens any APCF XML file.
<version>1.0</version>	The mandatory tag that specifies the APCF implementation version. Currently the only version is 1.0.
<application>...</application>	The mandatory tag that wraps the body of the XML description.
<id> text </id>	The mandatory tag that describes this particular APCF functionality.
<apcf-entities>...</apcf-entities>	The mandatory tag that wraps a single or multiple APCF entities.
<js-object>…</js-object> <html-object>…</html-object> <process-request-header>...</process-request-header> <process-response-header>...</process-response-header> <preprocess-response-body>...</preprocess-response-body> <postprocess-response-body>...</postprocess-response-body>	One of these tags specifies type of content or the stage at which the APCF processing should take place.

Table 91-1 APCF XML Tags (continued)

Tag	Use
<conditions>… </conditions>	A child element of the pre/post-process tags that specifies criteria for processing such as: http-version (such as 1.1, 1.0, 0.9) http-method (get, put, post, webdav) http-scheme (“http/”, “https/”, other) server-regexp regular expression containing ("a".."z" | "A".."Z" | "0".."9" | ".-_*[]?")) server-fnmatch (regular expression containing ("a".."z" | "A".."Z" | "0".."9" | ".-_*[]?+()\{},")), user-agent-regexp user-agent-fnmatch request-uri-regexp request-uri-fnmatch If more than one of condition tags is present, the ASA performs a logical AND for all tags.
<action> … </action>	Wraps one or more actions to perform on the content under specified conditions; you can use the following tags to define these actions (shown below): <do>, <sed-script>, <rewrite-header>, <add-header>, <delete-header>.
<do>…</do>	Child element of the action tag used to define one of the following actions: <no-rewrite/>—Do not mangle the content received from the remote server. <no-toolbar/>—Do not insert the toolbar. <no-gzip/>—Do not compress the content. <force-cache/>—Preserve the original caching instructions. <force-no-cache/>—Make object non-cacheable. < downgrade-http-version-on-backend>—Use HTTP/1.0 when sending the request to remote server.
<sed-script> TEXT </sed-script>	Child element of the action tag used to change the content of text-based objects. The Text must be a valid Sed script. The <sed-script> applies to the <conditions> tag defined before it.
<rewrite-header></rewrite-header>	Child element of the action tag. Changes the value of the HTTP header specified in the child element <header> tag shown below.
<add-header></add-header>	Child element of the action tag used to add a new HTTP header specified in the child element <header> tag shown below.

Table 91-1 APCF XML Tags (continued)

Tag	Use
<delete-header></delete-header>	Child element of the action tag used to delete the specified HTTP header specified by the child element <header> tag shown below.
<header></header>	Specifies the name HTTP header to be rewritten, added, or deleted. For example, the following tag changes the value of the HTTP header named Connection: <rewrite-header> <header>Connection</header> <value>close</value> </rewrite-header>

Configuration Examples for APCF

Example:

```
<APCF>
<version>1.0</version>
<application>
  <id>Do not compress content from example.com</id>
  <apcf-entities>
      <process-request-header>
         <conditions>
           <server-fnmatch>*.example.com</server-fnmatch>
         </conditions>
           <action>
             <do><no-gzip/></do>
           </action>
      </process-request-header>
  </apcf-entities>
</application>
</APCF>
```

Example:

```
<APCF>
<version>1.0</version>
<application>
 <id>Change MIME type for all .xyz objects</id>
 <apcf-entities>
      <process-response-header>
        <conditions>
            <request-uri-fnmatch>*.xyz</request-uri-fnmatch>
        </conditions>
         <action>
           <rewrite-header>
                <header>Content-Type</header>
                <value>text/html</value>
           </rewrite-header>
         </action>
      </process-response-header>
 </apcf-entities>
</application>
</APCF>
```

Managing Passwords

Optionally, you can configure the ASA to warn end users when their passwords are about to expire.

The ASA supports password management for the RADIUS and LDAP protocols. It supports the "password-expire-in-days" option for LDAP only.

You can configure password management for IPsec remote access and SSL VPN tunnel-groups.When you configure password management, the ASA notifies the remote user at login that the user's current password is about to expire or has expired. The ASA 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 command is valid for AAA servers that support such notification.

The ASA, 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

The RADIUS server (for example, Cisco ACS) could proxy the authentication request to another authentication server. However, from the ASA perspective, it is talking only to a RADIUS server.

Prerequisites

- 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.

  If you are using an LDAP directory server for authentication, password management is supported with the Sun Java System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory.

  Sun—The DN configured on the ASA 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.Restrictions
- Some RADIUS servers that support MSCHAP currently do not support MSCHAPv2. This command requires MSCHAPv2 so check with your vendor.
- Password management is *not* supported for any of these connection types for Kerberos/Active Directory (Windows password) or NT 4.0 Domain.
- For LDAP, the method to change a password is proprietary for the different LDAP servers on the market. Currently, the ASA implements the proprietary password management logic only for Microsoft Active Directory and Sun LDAP servers.
- The ASA ignores this command if RADIUS or LDAP authentication has not been configured.

Detailed Steps

Step 1 Navigate to Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles > Add or Edit > Advanced > General > Password Management.

Step 2 Click the Enable password management option.

Adding the Cisco Authentication Scheme to SiteMinder

In addition to configuring the ASA for SSO with SiteMinder, you must also configure your CA SiteMinder Policy Server with the Cisco authentication scheme, a Java plug-in you download from the Cisco web site.

Prerequisites

Configuring the SiteMinder Policy Server requires experience with SiteMinder.

Detailed Steps

This section presents general tasks, not a complete procedure. 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 on the ASA.

  You configure the secret on the ASA using the **policy-server-secret** command at the command line interface.
- 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/cisco/software/navigator.html and copy it to the default library directory for the SiteMinder server. This .jar file is also available on the Cisco ASA CD.

Configuring the SAML POST SSO Server

Use the SAML server documentation provided by the server software vendor to configure the SAML server in Relying Party mode.The following steps list the specific parameters required to configure the SAML Server for Browser Post Profile:

Detailed Steps

Step 1 Configure the SAML server parameters to represent the asserting party (the ASA):

- Recipient consumer 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>

Configuring SSO with the HTTP Form Protocol

This section describes using the HTTP Form protocol for SSO. HTTP Form protocol is an approach to SSO authentication that can also qualify as a AAA method. It provides a secure method for exchanging authentication information between users of clientless SSL VPN and authenticating web servers. You can use it in conjunction with other AAA servers such as RADIUS or LDAP servers.

Prerequisites

To configure SSO with the HTTP protocol correctly, you must have a thorough working knowledge of authentication and HTTP protocol exchanges.

Restrictions

As a common protocol, it is applicable only when the following conditions are met for the web server application used for authentication:

- The authentication cookie must be set for successful request and not set for unauthorized logons. In this case, ASA cannot distinguish successful from failed authentication.

Detailed Steps

The ASA again serves as a proxy for users of clientless SSL VPN to an authenticating web server but, in this case, it uses HTTP Form protocol and the POST method for requests. You must configure the ASA to send and receive form data. Figure 91-4 illustrates the following SSO authentication steps:

Step 1 A user of clientless SSL VPN first enters a username and password to log into the clientless SSL VPN server on the ASA.

Step 2 The clientless SSL VPN server acts as a proxy for the user and forwards the form data (username and password) to an authenticating web server using a POST authentication request.

Step 3 If the authenticating web server approves the user data, it returns an authentication cookie to the clientless SSL VPN server where it is stored on behalf of the user.

Step 4 The clientless SSL VPN server establishes a tunnel to the user.

Step 5 The user can now access other websites within the protected SSO environment without reentering a username and password.

Figure 91-4 SSO Authentication Using HTTP Forms

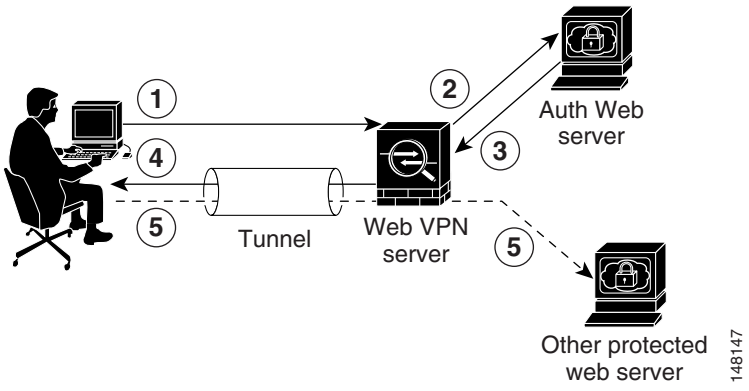

While you would expect to configure form parameters that let the ASA include POST data such as the username and password, you initially might not be aware of additional hidden parameters that the web server requires. Some authentication applications expect hidden data which is neither visible to nor entered by the user. You can, however, discover hidden parameters the authenticating web server expects by making a direct authentication request to the web server from your browser without the ASA in the middle acting as a proxy. Analyzing the web server response using an HTTP header analyzer reveals hidden parameters in a format similar to the following:

```
<param name>=<URL encoded value>&<param name>=<URL encoded>
```

Some hidden parameters are mandatory and some are optional. If the web server requires data for a hidden parameter, it rejects any authentication POST request that omits that data. Because a header analyzer does not tell you if a hidden parameter is mandatory or not, we recommend that you include all hidden parameters until you determine which are mandatory.

Gathering HTTP Form Data

This section presents the steps for discovering and gathering necessary HTTP Form data. If you do not know what parameters the authenticating web server requires, you can gather parameter data by analyzing an authentication exchange using the following steps:

Prerequisites

These steps require a browser and an HTTP header analyzer.

Detailed Steps

Step 1 Start your browser and HTTP header analyzer, and connect directly to the web server login page without going through the ASA.

Step 2 After the web server login page has loaded in your browser, examine the login sequence to determine if a cookie is being set during the exchange. If the web server has loaded a cookie with the login page, configure this login page URL as the *start-URL*.

Step 3 Enter the username and password to log in to the web server, and press **Enter**. This action generates the authentication POST request that you examine using the HTTP header analyzer.

An example POST request—with host HTTP header and body—follows:

```
POST
/emco/myemco/authc/forms/MCOlogin.fcc?TYPE=33554433&REALMOID=06-000430e1-7443-125c-ac05
-83846dc90034&GUID=&SMAUTHREASON=0&METHOD=GET&SMAGENTNAME=$SM$5FZmjnk3DRNwNjk2KcqVCFbIr
NT9%2bJ0H0KPshFtg6rB1UV2PxkHqLw%3d%3d&TARGET=https%3A%2F%2Fwww.example.com%2Femco%2Fmye
mco%2FHTTP/1.1

Host: www.example.com

(BODY)

SMENC=ISO-8859-1&SMLOCALE=US-EN&USERID=Anyuser&USER_PASSWORD=XXXXXX&target=https%3A%2F%
2Fwww.example.com%2Femco%2Fmyemco%2F&smauthreason=0
```

Step 4 Examine the POST request and copy the protocol, host, and the complete URL to configure the action-uri parameter.

Step 5 Examine the POST request body and copy the following:

a. Username parameter. In the preceding example, this parameter is *USERID*, not the value *anyuser*.

b. Password parameter. In the preceding example, this parameter is *USER_PASSWORD*.

c. Hidden parameter. This parameter is everything in the POST body except the username and password parameters. In the preceding example, the hidden parameter is: SMENC=ISO-8859-1&SMLOCALE=US-EN&target=https%3A%2F%2Fwww.example.com%2Femco%2Fmyemco%2F&smauthreason=0

Figure 91-5 highlights the action URI, hidden, username and password parameters within sample output from an HTTP analyzer. This is only an example; output varies widely across different websites.

Figure 91-5 Action-uri, hidden, username and password parameters

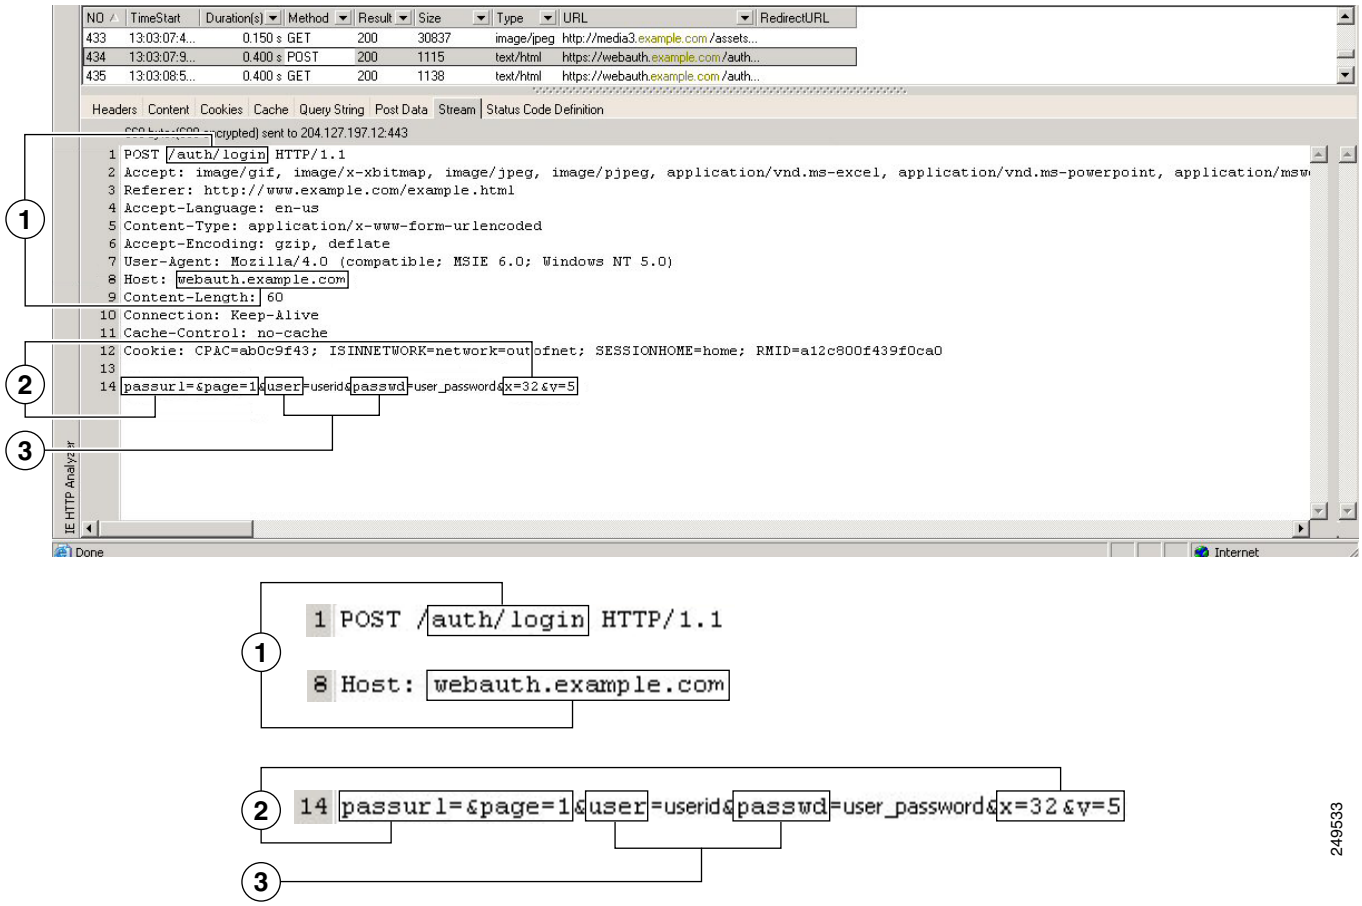

Step 6 If you successfully log in to the web server, examine the server response with the HTTP header analyzer to locate the name of the session cookie set by the server in your browser. This is the **auth-cookie-name** parameter.

In the following server response header, the name of the session cookie is SMSESSION. You just need the name, not the value.Figure 91-6 shows an example of authorization cookies in HTTP analyzer output. This is only an example; output varies widely across different websites.

Figure 91-6 Authorization cookies in sample HTTP analyzer output

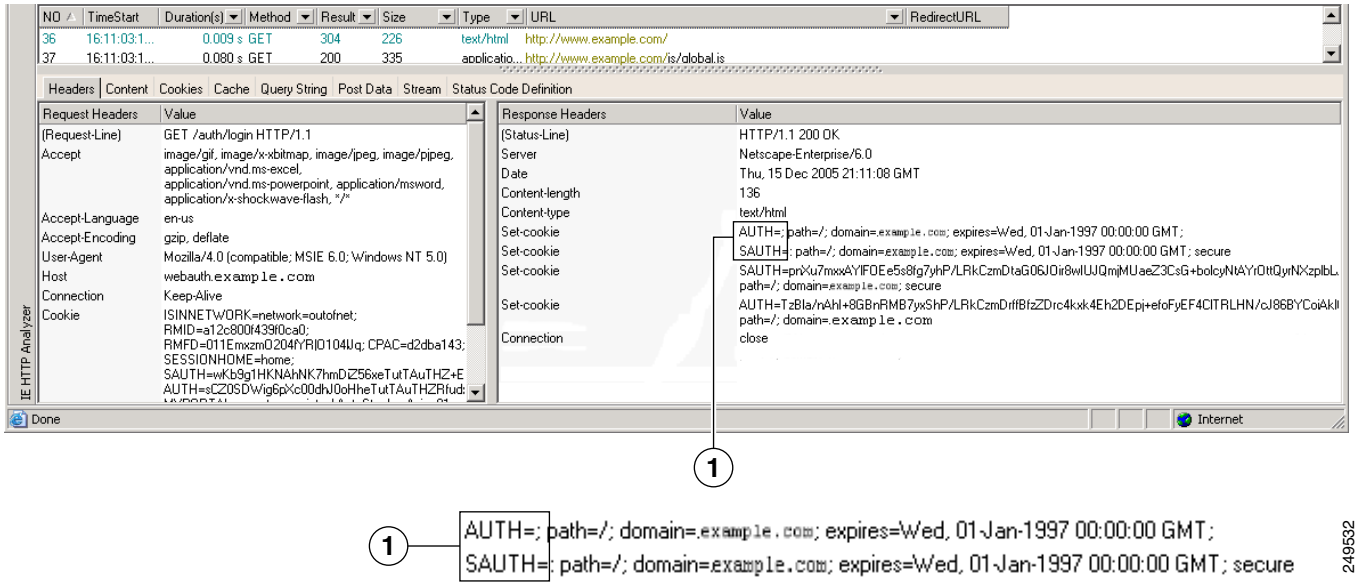

1	Authorization cookies

Step 7 In some cases, the server may set the same cookie regardless of whether the authentication was successful or not, and such a cookie is unacceptable for SSO purposes. To confirm that the cookies are different, repeat Step 1 through Step 6 using invalid login credentials and then compare the “failure” cookie with the “success” cookie.

You now have the necessary parameter data to configure the ASA for SSO with HTTP Form protocol.

Using 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 ASA passes the login credentials that the user of clientless SSL VPN entered to log in to the ASA (username and password) to those particular internal servers. You configure the ASA to respond to a specific authentication method for a particular range of servers. The authentication methods you can configure the ASA to respond to consists of authentication using Basic (HTTP), NTLM, FTP and CIFS, or all of these methods.

If the lookup of the username and password fails on the ASA, an empty string is substituted, and the behavior converts back as if no auto sign-on is available.

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 ASA to support this solution, see the "SSO Servers" section on page 91-61.

The following fields are displayed:

- IP Address—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—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—Displays a URI mask that identifies the servers configured with the Add/Edit Auto Signon dialog box.
- Authentication Type—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.

Restrictions

- Do not enable auto signon for servers that do not require authentication or that use credentials different from the ASA. When auto signon is enabled, the ASA passes on the login credentials that the user entered to log into the ASA regardless of what credentials are in user storage.
- 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 ASA does not pass the user login credentials to that server.

Detailed Steps

Step 1 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.

Step 2 Click to delete an auto signon instruction selected in the Auto Signon table.

Step 3 Click **IP Block** 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.

Step 4 Click URI to specify a server supporting auto signon by URI, then enter the URI in the field next to this button.

Step 5 Determine the authentication method assigned to the servers. For the specified range of servers, the ASA 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.

Note NTLM is not the most secure choice and other better alternatives exist. NTLM does not support recent cryptographic methods such as AES or SHA-256. Cisco recommends that applications not use NTLM.

- 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.

Accessing Virtual Desktop Infrastructure (VDI)

In a VDI model, administrators publish enterprise applications or desktops pre-loaded with enterprise applications, and end users remotely access these applications. These virtualized resources appear just as any other resources, such as email, so that users do not need to go through a Citrix Access Gateway to access them. Users log onto the ASA using Citrix Receiver mobile client, and the ASA connects to a pre-defined Citrix XenApp or XenDesktop Server. The administrator must configure the Citrix server's address and logon credentials under Group Policy so that when users connect to their Citrix Virtualized resource, they enter the ASA's SSL VPN IP address and credentials instead of pointing to the Citrix Server's address and credentials. When the ASA has verified the credentials, the receiver client starts to retrieve entitled applications through the ASA.

Supported Mobile Devices

- iPad—Citrix Receiver version 4.x or later
- iPhone/iTouch—Citrix Receiver version 4.x or later
- Android 2.x/3.x/4.0/4.1 phone—Citrix Receiver version 2.x or later
- Android 4.0 phone—Citrix Receiver version 2.x or later

Limitations

Citrix Receiver clients access only one XenApp/XenDesktop server at a time.

Detailed Steps

Step 1 Browse to Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies.

Step 2 Edit the DfltGrpPolicy and expand the More options menu from the left-side menu.

Step 3 Choose **VDI Access**. Click **Add** or **Edit** to provide VDI server details.

- Server (Host Name or IP Address)—Address of the XenApp or XenDesktop server. This value can be a clientless macro.
- Port Number (Optional)—Port number for connecting to the Citrix server. This value can be a clientless macro.
- Active Directory Domain Name—Domain for logging into the virtualization infrastructure server. This value can be a clientless macro.
- Use SSL Connection—Check the checkbox if you want the server to connect using SSL.
- Username—Username for logging into the virtualization infrastructure server. This value can be a clientless macro.
- Password—Password for logging into the virtualization infrastructure server. This value can be a clientless macro.

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.

Detailed Steps

Step 1 Click none or choose the file server protocol (smb or ftp) from the User Storage Location 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 ASA uses an internal algorithm to store the data in an encrypted form to safeguard it.

Step 2 Type the string, if required, for the security appliance to pass to provide user access to the storage location.

Step 3 Choose one of the following options from the Storage Objects drop-down menu to specify the objects the server uses in association with the user. The ASA store these objects to support clientless SSL VPN connections.

- cookies,credentials
- cookies
- credentials

Step 4 Enter the limit in KB transaction size 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.

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 the configured certificate that you want to employ in Java object signing from the drop down list.

To configure a Java Code Signer, choose Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > 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**.

Encoding

With encoding, you can 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 ASA 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.

Detailed Steps

Step 1 Global Encoding Type 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

- unicode
- windows-1252
- none

Note 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 ASA configuration.

Step 2 Enter the name or IP address of a CIFS server for which the encoding requirement differs from the "Global Encoding Type" attribute setting. The ASA retains the case you specify, although it ignores the case when matching the name to a server.

Step 3 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 ASA configuration.

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.

Detailed Steps

Step 1 Select **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Content Cache**.

Step 2 If **Enable Cache** is unchecked, check it.

Step 3 Define the terms for caching.

- Maximum Object Size—Enter the maximum size in KB of a document that the ASA can cache. The ASA 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 ASA can cache. The ASA 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 ASA 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 ASA to cache objects that have neither a last-modified time stamp nor an explicit server-set expiry time.
- Cache static content—Check 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.

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 ASA. The ASA therefore lets you create rewrite rules that let users browse certain sites and applications without going through the ASA. This is similar to split-tunneling in a 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.

"Configuration Example for Content Rewrite Rules" shows example content rewrite rules.

Note These improvements were made to Content Rewriter in ASA 9.0:

- Content rewrite added support for HTML5.

- The clientless SSL VPN rewriter engines were significantly improved to provide better quality and efficacy. As a result, you can expect a better end-user experience for clientless SSL VPN users.

Detailed Steps

The Content Rewrite table has the following columns:

- 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.

The following steps explain how to add a rewrite entry or edit a selected rewrite entry.

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Content Rewrite**.

Step 2 Click Add or Edit to create or update an content rewriting rule.

Step 3 Enable content rewrite must be checked to enable this rule.

Step 4 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.

Step 5 (Optional) Provide an alphanumeric string that describes the rule, maximum 128 characters.

Step 6 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.

Configuration Example for Content Rewrite Rules

Table 91-2 Content Rewrite Rules

Function	Enable content rewrite	Rule Number	Rule Name	Resource Mask
Disable rewriter for HTTP URLs at youtube.com	Unchecked	1	no-rewrite-youtube	*.youtube.com/*
Enable rewriter for all HTTP URLs that do not match above rules	Check	65,535	rewrite-all	*

Configuring Browser Access to Plug-ins

The following sections describe the integration of browser plug-ins for clientless SSL VPN browser access:

- Preparing the Security Appliance for a Plug-in, page 91-34
- Installing Plug-ins Redistributed By Cisco, page 91-34
- Providing Access to a Citrix XenApp Server, page 91-36

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 ASA 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.

The ASA 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 ASA 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 91-3 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.

* Not a recommended plug-in.

Table 91-3 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 MetaFrame Services	ica://
rdp	Terminal Servers	rdp://
rdp2*	Terminal Servers Vista	rdp2://
ssh,telnet	Secure Shell	ssh://
	Telnet services (supporting v1 and v2)	telnet://
vnc	Virtual Network Computing services	vnc://

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.The plug-ins support single sign-on (SSO). Refer to the "Configuring SSO with the HTTP Form Protocol" section on page 91-22 for implementation details.

The minimum access rights required for remote use belong to the guest privilege mode.

Prerequisites

- Clientless SSL VPN must be enabled on the ASA to provide remote access to the plug-ins.
- To configure SSO support for a plug-in, you install the plug-in, add a bookmark entry to display a link to the server, and specify SSO support when adding the bookmark.
- The minimum access rights required for remote use belong to the guest privilege mode.

- Plug-ins require ActiveX or Oracle Java Runtime Environment (JRE), see the compatibility matrix for version requirements.

Restrictions

Note The remote desktop protocol plug-in does not support load balancing with a session broker. Because of the way the protocol handles the redirect from the session broker, the connection fails. If a session broker is not used, the plug-in works.

- The plug-ins support single sign-on (SSO). They use the *same* credentials entered to open the clientless SSL VPN session. Because the plug-ins do not support macro substitution, you do not have the options to perform SSO on different fields such as the internal domain password or on an attribute on a RADIUS or LDAP server.
- A stateful failover does not retain sessions established using plug-ins. Users must reconnect following a failover.
- If you use stateless failover instead of stateful failover, clientless features such as bookmarks, customization, and dynamic access-policies are not synchronized between the failover ASA pairs. In the event of a failover, these features do not work.

Preparing the Security Appliance for a Plug-in

Before installing a plug-in, prepare the ASA as follows:

Prerequisites

Make sure clientless SSL VPN (“webvpn”) is enabled on an ASA interface.

Restrictions

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 ASA. The remote PC must be able to use DNS or an entry in the System32\drivers\etc\hosts file to resolve the FQDN.

Go to the section that identifies the type of plug-in you want to provide for clientless SSL VPN access.

- Installing Plug-ins Redistributed By Cisco, page 91-34
- Providing Access to a Citrix XenApp Server, page 91-36

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.

Prerequisites

Make sure clientless SSL VPN (“webvpn”) is enabled on an interface on the ASA. To do so, enter the **show running-config** command.

Plug-ins Redistributed by Cisco

Table 91-4

Protocol	Description	Source of Redistributed Plug-in *
RDP	HOBLink JWT is a Native Java RDP client which supports RDP 7.0 for remote access to Windows Terminal services. **Note** Since properoRDP and HOBSoft both use RDP as the plugin protocol, customers will not be able to use HOBSoft for one set of users and properoRDP for the rest via GroupPolicy etc. Use this plug-in if you want to use a single plug-in for Windows, Mac OS X and Linux operating systems. See This Document for a list of all supported operating systems.	Customers can download this plugin from these locations: HOBSoft Cisco.com Proper Java RDP
RDP	Accesses Microsoft Terminal Services hosted by Windows Vista and Windows 2003 R2. Supports Remote Desktop ActiveX Control. 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.	The original source of the redistributed plug-in is http://properjavardp.sourceforge.net/
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. We do not recommend using this plug-in, instead, use the RDP plug-in above.	The original source of the redistributed plug-in is http://properjavardp.sourceforge.net/
SSH	The Secure Shell-Telnet plug-in lets the remote user establish a Secure Shell (v1 or v2) 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.)	The web site containing the source of the redistributed plug-in is http://javassh.org/
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 (also known as VNC server or service) turned on. This version changes the default color of the text and contains updated French and Japanese help files.	The web site containing the source of the redistributed plug-in is http://www.tightvnc.com/

```
* Consult the plug-in documentation for information on deployment configuration and
restrictions.
```

These plug-ins are available on the Cisco Adaptive Security Appliance Software Download site.

Detailed Steps

Follow these steps to provide clientless SSL VPN browser access to a plug-in redistributed by Cisco.

Step 1 Create a temporary directory named plugins on the computer you use to establish ASDM sessions with the ASA, and download the plug-ins you want from the Cisco website to the plugins directory.

Step 2 Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Client-Server Plug-ins**.

This pane displays the currently loaded plug-ins that are available to clientless SSL sessions. The hash and date of these plug-ins are also provided.

Step 3 Click **Import**.

The Import Client-Server Plug-in dialog box opens.

Step 4 Use the following descriptions to enter the Import Client-Server Plug-in dialog box field values.

- Plug-in Name—Select one of the following values:
  - **ica** to provide plug-in access to Citrix MetaFrame or Web Interface services
  - **rdp** to provide plug-in access to Remote Desktop Protocol services
  - **ssh,telnet** to provide plug-in access to *both* Secure Shell and Telnet services
  - **vnc** to provide plug-in access to Virtual Network Computing services

Note Any undocumented options in this menu are experimental and are not supported.

- Select the location of the plugin file—Select one of the following options and insert a path into its text field.
  - Local computer—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—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—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 5 Click **Import Now**.

Step 6 Click **Apply**.

The plug-in is now available for future clientless SSL VPN sessions.

Providing Access to a Citrix XenApp Server

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 XenApp Server Client.

With a Citrix plug-in installed on the ASA, clientless SSL VPN users can use a connection to the ASA to access Citrix XenApp 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 XenApp Server for Clientless SSL VPN Access

- Creating and Installing the Citrix Plug-in

Preparing the Citrix XenApp Server for Clientless SSL VPN Access

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 XenApp Server.

Note If you are not already providing support for a plug-in, you must follow the instructions in the"Preparing the Security Appliance for a Plug-in" section on page 91-34 before using this section.

Creating and Installing the Citrix Plug-in

To create and install the Citrix plug-in, perform the following steps:

Detailed Steps

Step 1 Download theICA plug-in file from the Cisco Software Download web site.

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.

On Citrix download site, select Citrix Receiver, Receivers by Platform, and click Find. Expand Receiver for Other Platforms, and download Receiver for Java. JICAComponents.tar.gz a gzip'd tar file, which you can open with 7-Zip or other unix-compatible tools.

Step 3 Extract the following files from the Citrix Java client, then add them to the ica-plugin.zip file:

- JICA-configN.jar
- JICAEngN.jar

You can use WinZip to perform this step.

Step 4 Ensure the EULA included with the Citrix Java client grants you the rights and permissions to deploy the client on your web servers.

Step 5 Install the plug-in by using ASDM, or entering the following CLI command in privileged EXEC mode:

import webvpn plug-in protocol ica *URL*

URL is the host name or IP address and path to the ica-plugin.zip file.

Note We recommend that you add a bookmark to make it easy for users to connect. Adding a bookmark is required if you want to provide SSO support for Citrix sessions.
We also recommend that you use URL parameters in the bookmark the provide convenient viewing, for example:
ica://10.56.1.114/?DesiredColor=4&DesiredHRes=1024&DesiredVRes=768

Step 6 Establish an SSL VPN clientless session and click the bookmark or enter the URL for the Citrix server.

Use the Client for Java Administrator's Guide as needed.

Microsoft Kerberos Constrained Delegation Solution

Many organizations want to authenticate their Clientless VPN users and extend their authentication credentials seamlessly to web-based resources using authentication methods beyond what the ASA SSO feature can offer today. With the growing demand to authenticate remote access users with Smart Cards and One-time Passwords (OTP), the SSO feature falls short in meeting that demand, because it only forwards conventional user credentials, such as static username and password, to clientless web-based resources when authentication is required.

For example, neither certificate- or OTP-based authentication methods encompass a conventional username and password necessary for the ASA to seamlessly perform SSO access to web-based resources. When authenticating with a certificate, a username and password is not required for the ASA to extend to web-based resources, making it an unsupported authentication method for SSO. On the other hand, OTP does include a static username; however, the password is dynamic and will subsequently change throughout the VPN session. In general, Web-based resources are configured to accept static usernames and passwords, thus also making OTP an unsupported authentication method for SSO.

Microsoft's Kerberos Constrained Delegation (KCD), a new feature introduced in software release 8.4 of the ASA, provides access to Kerberos-protected Web applications in the private network. With this benefit, you can seamlessly extend certificate- and OTP-based authentication methods to web applications. Thus, with SSO and KCD working together although independently, many organizations can now authenticate their clientless VPN users and extend their authentication credentials seamlessly to web applications using all authentication methods supported by the ASA.

Requirements

In order for the **kcd-server** command to function, the ASA must establish a trust relationship between the *source* domain (the domain where the ASA resides) and the *target* or *resource* domain (the domain where the web services reside). The ASA, using its unique format, crosses the certification path from the source to the destination domain and acquires the necessary tickets on behalf of the remote access user to access the services.

This crossing of the certificate path is called cross-realm authentication. During each phase of cross-realm authentication, the ASA relies on the credentials at a particular domain and the trust relationship with the subsequent domain.

Understanding How KCD Works

Kerberos relies on a trusted third party to validate the digital identity of entities in a network. These entities (such as users, host machines, and services running on hosts) are called principals and must be present in the same domain. Instead of secret keys, Kerberos uses tickets to authenticate a client to a server. The ticket is derived from the secret key and consists of the client's identity, an encrypted session key, and flags. Each ticket is issued by the key distribution center and has a set lifetime.

The Kerberos security system is a network authentication protocol used to authenticate entities (users, computers, or applications) and protect network transmissions by scrambling the data so that only the device that the information was intended for can decrypt it. You can configure KCD to provide Clientless SSL VPN (also known as WebVPN) users with SSO access to any web services protected by Kerberos. Examples of such web services or applications include Outlook Web Access (OWA), Sharepoint, and Internet Information Server (IIS).

Two extensions to the Kerberos protocol were implemented: *protocol transition* and *constrained delegation*. These extensions allow the Clientless or WebVPN remote access users to access Kerberos authenticated applications in the private network.

The *protocol transition* provides you with increased flexibility and security by supporting different authentication mechanisms at the user authentication level and by switching to the Kerberos protocol for security features (such as mutual authentication and constrained delegation) in subsequent application layers. *Constrained delegation* provides a way for domain administrators to specify and enforce application trust boundaries by limiting where application services can act on a user's behalf. This flexibility improves application security designs by reducing the chance of compromise by an untrusted service.

For more information on constrained delegation, see RFC 1510 via the IETF website (http://www.ietf.org).

Authentication Flow with KCD

Figure 91-7 depicts the packet and process flow a user will experience directly and indirectly when accessing resources trusted for delegation via the clientless portal. This process assumes that the following tasks have been completed:

- Configured KCD on ASA
- Joined the Windows Active Directory and ensured services are trusted for delegation
- Delegated ASA as a member of the Windows Active Directory domain

Figure 91-7 KCD Process

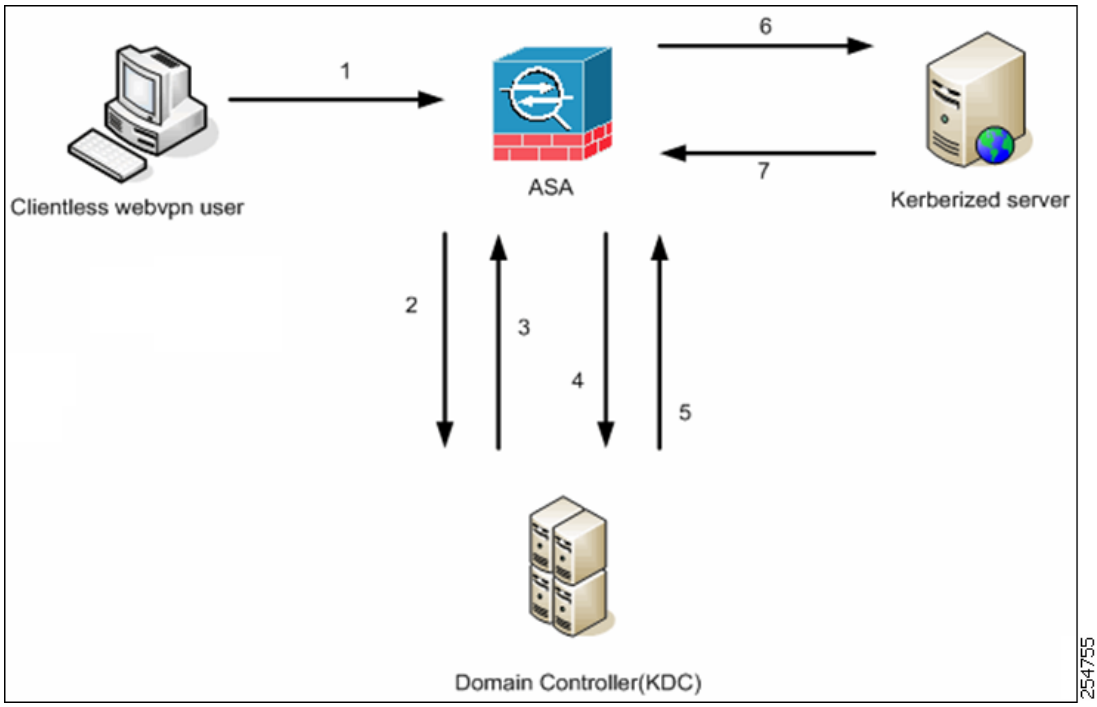

Note A clientless user session is authenticated by the ASA using the authentication mechanism configured for the user. (In the case of Smartcard credentials, ASA performs LDAP authorization with the userPrincipalName from the digital certificate against the Windows Active Directory).

1. After successful authentication, the user logs in to the ASA clientless portal page. The user accesses a Web service by entering a URL in the portal page or by clicking on the bookmark. If the Web service requires authentication, the server challenges ASA for credentials and sends a list of authentication methods supported by the server.

   Note KCD for Clientless SSL VPN is supported for all authentication methods (RADIUS, RSA/SDI, LDAP, digital certificates, and so on). Refer to the AAA Support table at http://www.cisco.com/en/US/partner/docs/security/asa/asa84/configuration/guide/access_aaa.html#wp1069492.

2. Based on the HTTP headers in the challenge, ASA determines whether the server requires Kerberos authentication. (This is part of the SPNEGO mechanism.) If connecting to a backend server requires Kerberos authentication, the ASA requests a service ticket for itself on behalf of the user from the key distribution center.
3. The key distribution center returns the requested tickets to the ASA. Even though these tickets are passed to the ASA, they contain the user's authorization data.ASA requests a service ticket from the KDC for the specific service that the user wants to access.

   Note Steps 1 to 3 comprise protocol transition. After these steps, any user who authenticates to ASA using a non-Kerberos authentication protocol is transparently authenticated to the key distribution center using Kerberos.

4. ASA requests a service ticket from the key distribution center for the specific service that the user wants to access.
5. The key distribution center returns a service ticket for the specific service to the ASA.
6. ASA uses the service ticket to request access to the web service.
7. The Web server authenticates the Kerberos service ticket and grants access to the service. The appropriate error message is displayed and requires acknowledgement if there is an authentication failure. If the Kerberos authentication fails, the expected behavior is to fall back to basic authentication.

Adding Windows Service Account in Active Directory

The KCD implementation on the ASA requires a service account, or in other words, an Active Directory user account with privileges necessary to add computers, such as adding the ASA to the domain. For our example, the Active Directory username JohnDoe depicts a service account with the required privileges. For more information on how to implement user privileges in Active Directory, contact Microsoft Support or visit http://microsoft.com.

Configuring DNS for KCD

This section outlines configuration procedures necessary to configure DNS on the ASA. When using KCD as the authentication delegation method on the ASA, DNS is required to enable hostname resolution and communication between the ASA, Domain Controller (DC), and services trusted for delegation.

Step 1 From ASDM, navigate to **Configuration > Remote Access VPN > DNS** and configure the DNS setup as shown in Figure 91-8:

- DNS Server Group—Enter the DNS server IP address(es), such as 192.168.0.3.
- Domain Name—Enter the domain name in which the DC is a member.

Step 2 Enable DNS Lookup on the appropriate interface. Clientless VPN deployments require DNS Lookups via the internal corporate network, typically the *inside* interface.

Figure 91-8 ASA DNS Configuration Example

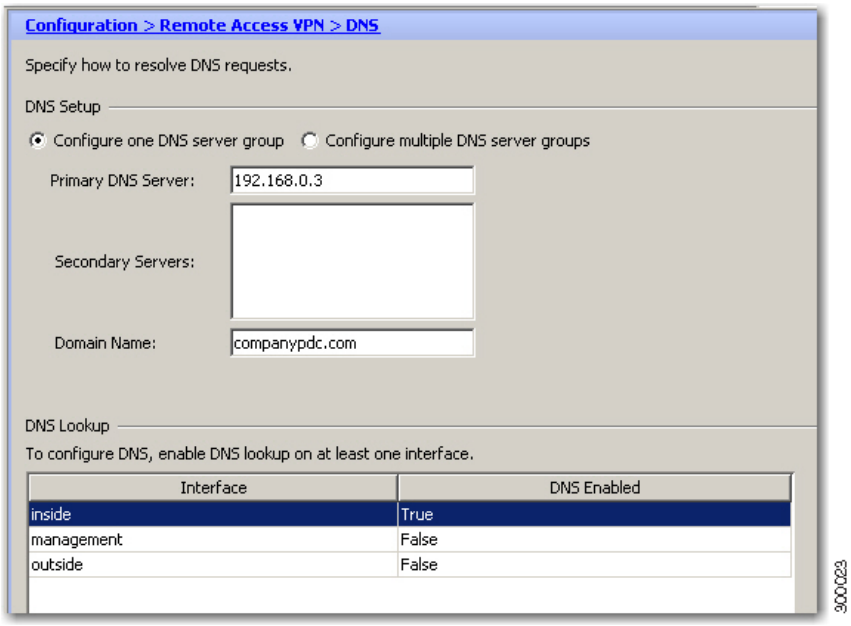

Configuring the ASA to Join the Active Directory Domain

This section outlines configuration procedures necessary to enable the ASA to act as part of the Active Directory domain. KCD requires the ASA to be a member of the Active Directory domain. This configuration enables the functionality necessary for constrained delegation transactions between the ASA and the KCD server.

Step 1 From ASDM, navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Microsoft KCD Server**, as shown in Figure 91-9.

Step 2 Click **New** to add a Kerberos Server Group for Constrained Delegation and configure the following (see Figure 91-9):

- Server Group Configuration

  - Server Group Name—Define the name of the constrained delegation configuration on the ASA, such as MSKCD, which is the default value. You can configure multiple server groups for redundancy; however, you can only assign one server group to the KCD server configuration used to request service tickets on behalf of VPN users.
  - Reactivation Mode—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. Depletion is the default configuration.
  - Dead Time—If you choose the Depletion reactivation mode, you must add a dead time interval. Ten minutes is the default configuration. The interval represents the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers.
  - Max Failed Attempts—Set the number of failed connection attempts allowed before declaring an unresponsive server to be inactive. Three attempts is the default.
- Server Configuration
  - Interface Name—Choose the interface on which the server resides. In general, authentication server deployments reside on the internal corporate network, typically via the *inside* interface.
  - Server Name—Define the hostname of the domain controller, such as ServerHostName.
  - Timeout—Specify the maximum time, in seconds, to wait for a response from the server. Ten seconds is the default.
- Kerberos Parameter
  - Server Port—88 is the default and the standard port used for KCD.
  - Retry Interval—Choose the desired retry interval. Ten seconds is the default configuration.
  - Realm—Enter the domain name of the DC in all uppercase. The KCD configuration on the ASA requires the realm value to be in uppercase. A realm is an authentication domain. A service can accept authentication credentials only from entities in the same realm. The realm must match the domain name which the ASA joins.

Figure 91-9 KCD Server Group Configuration

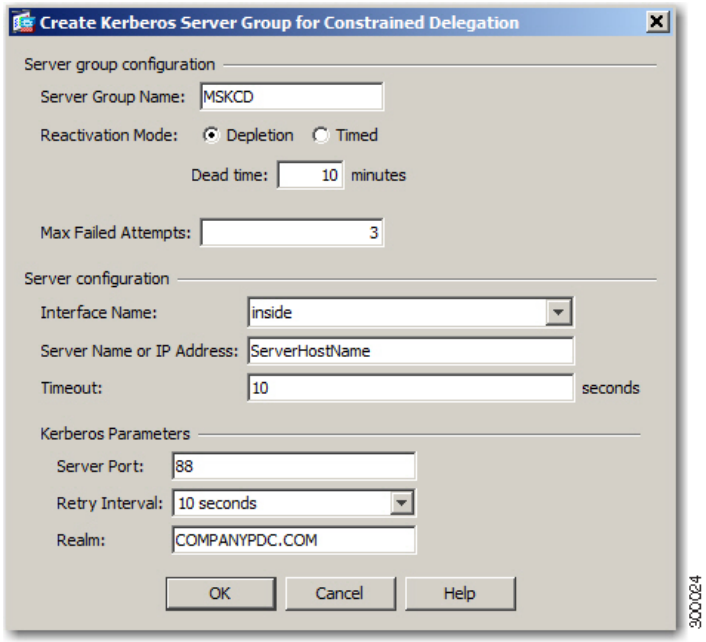

Step 3 Click **OK** to apply your configuration and then configure the Microsoft KCD Server to request service tickets on behalf of the remote access user (see Figure 91-9). The Microsoft KCD Server configuration window appears upon clicking **OK**.

Configuring Kerberos Server Groups

The Kerberos Server Group for Constrained Delegation, MSKCD, is automatically applied to the KCD Server Configuration. You can also configure Kerberos Server groups and manage them under **Configuration > Remote Access VPN > AAA/Local User > AAA Server Groups**.

Step 1 Under the Server Access Credential section, configure the following:

- Username—Define a Service Account (Active Directory username) such as JohnDoe, which has been granted privileges necessary to add computer accounts to the Active Directory domain. The username does not correspond to a specific administrative user but simply a user with service-level privileges. This service account is used by the ASA to add a computer account for itself to the Active Directory domain at every reboot. You must configure the computer account separately to request Kerberos tickets on behalf of the remote users.

  Note Administrative privileges are required for initial join. A user with service-level priviledges on the domain controller will not get access.

- Password—Define the password associated with the username (such as Cisco123). The password does not correspond to a specific password but simply a service-level password privilege to add a device on the Window domain controller.

Step 2 Under the Server Group Configuration section, configure the following:

- Reactivation Mode—Click 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. Depletion is the default configuration.
- Dead Time—If you choose the Depletion reactivation mode, you must add a dead time interval. The interval represents the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers. Ten minutes is the default.
- Max Failed Attempts—Set the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive. Three attempts is the default.

Note Under the Server Table section, the previously configured DC hostname, ServerHostName, was automatically applied to the KCD Server configuration (see Figure 91-10).

Figure 91-10 KCD Server Configuration

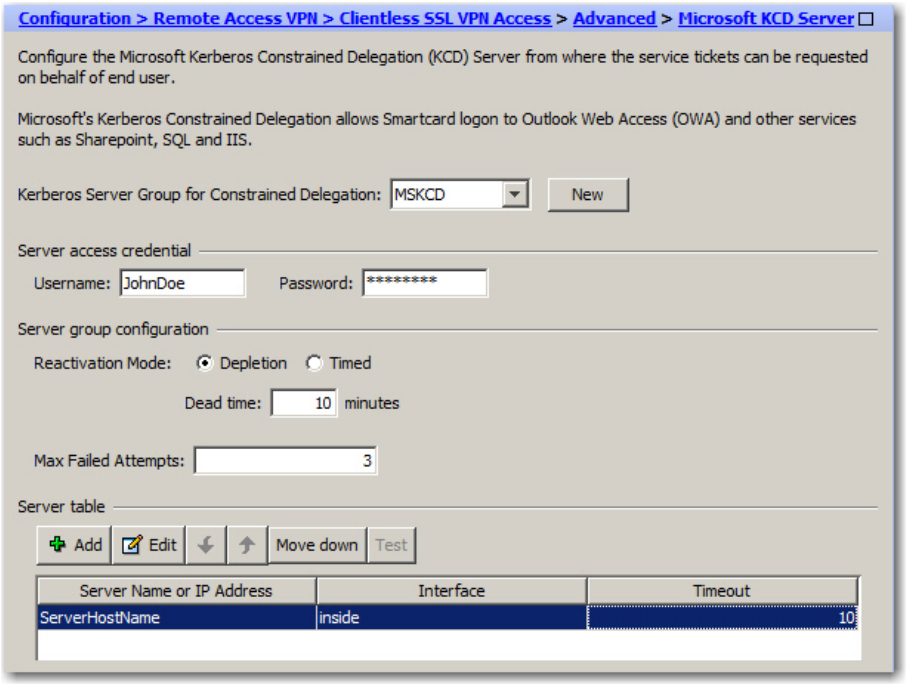

Step 3 Click **Apply**.

Note After applying your configuration, the ASA automatically starts the process of joining the Active Directory domain. The ASA's hostname appears in the Computers directory in Active Directory Users and Computers.

To confirm if the ASA has successfully joined the domain, execute the following command from the ASA prompt as shown in Figure 91-11:

show webvpn kcd

Figure 91-11 ASA Domain Membership Confirmation

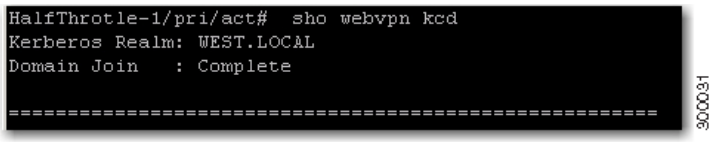

Configuring Bookmarks to Access the Kerberos Authenticated Services

To access Kerberos Authenticated Services such as Outlook Web Access using the ASA clientless portal, you must configure bookmark lists. Bookmark Lists are assigned and displayed to remote access users based on the VPN security policies they are associated with.

Restriction

When creating a bookmark to an application that uses Kerberos constrained delegation (KCD), do not check Enable Smart Tunnel.

Detailed Steps

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless VPN Access > Portal > Bookmarks** on the ASDM GUI.

Step 2 In Bookmark List, enter the URL to reference for the service location.

Configuring Application Access

The following sections describe how to enable smart tunnel access and port forwarding on clientless SSL VPN sessions, specify the applications to be provided with such access, and provide notes on using it:

- Configuring Smart Tunnel Access
- Configuring Smart Tunnel Log Off

Configuring Smart Tunnel Access

To configure smart tunnel access, you create a smart tunnel list containing one or more applications eligible for smart tunnel access, and the endpoint operating system associated with the list. 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. After creating a list, you assign it to one or more group policies or local user policies.

The following sections describe smart tunnels and how to configure them:

- About Smart Tunnels
- Why Smart Tunnels?
- Configuring a Smart Tunnel (Lotus example)

- Simplifying Configuration of Which Applications to Tunnel
- About Smart Tunnel Lists
- Creating a Smart Tunnel Auto Sign-On Server List
- Adding Servers to a Smart Tunnel Auto Sign-on Server List
- Enabling and Disabling Smart Tunnel Access

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 ASA 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 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 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.

Why Smart Tunnels?

Smart tunnel access lets a client TCP-based application use a browser-based VPN connection to access 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.

Prerequisites

See the *Supported VPN Platforms, Cisco ASA 5500 Series* for the platforms and browsers supported by ASA Release 9.0 smart tunnels.

The following requirements and limitations apply to smart tunnel access on Windows:

- ActiveX or Oracle Java Runtime Environment (JRE) 4 update 15 or later (JRE 6 or later recommended) on Windows must be enabled on the browser.
- Only Winsock 2, TCP-based applications are eligible for smart tunnel access.

- For Mac OS X only, Java Web Start must be enabled on the browser.

Restrictions

- Smart tunnel supports only proxies placed between computers running Microsoft Windows and the security appliance. Smart Tunnel uses the Internet Explorer configuration, which sets system-wide parameters in Windows. That configuration may include proxy information:
  - If a Windows computer requires a proxy to access the ASA, then there must be a static proxy entry in the client's browser, and the host to connect to must be in the client's list of proxy exceptions.
  - If a Windows computer does not require a proxy to access the ASA, but does require a proxy to access a host application, then the ASA must be in the client's list of proxy exceptions.

  Proxy systems can be defined the client's configuration of static proxy entry or automatic configuration, or by a PAC file. Only static proxy configurations are currently supported by Smart Tunnels.
- Kerberos constrained delegation (KCD) is not supported for smart tunnels.
- 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.
- With HTTP-based remote access, some subnets may block user access to the VPN gateway. To fix this, place a proxy in front of the ASA to route traffic between the web and the end user. That proxy must support the CONNECT method. For proxies that require authentication, Smart Tunnel supports only the basic digest authentication type.
- When smart tunnel starts, the ASA by default passes all browser traffic through the VPN session if the browser process is the same. The ASA only also does this if a tunnel-all policy (the default) 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.
- The Mac version of smart tunnel does not support POST bookmarks, form-based auto sign-on, or POST macro substitution.
- For Mac OS X users, only those 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.
- In Mac OS X, applications using TCP that are dynamically linked to the SSL library can work over a smart tunnel.
- Smart tunnel does not support the following on Mac OS X:
  - 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.

- Mac OS X 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).

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.

Detailed Steps

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 91-5 shows example Application ID strings and the associated paths required to support Lotus.

Table 91-5 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.

Simplifying Configuration of Which Applications to Tunnel

A smart tunnel application list is essentially a filter of what applications are granted access to the tunnel. The default is to allow access for all processes started by the browser. With Smart Tunnel enabled bookmark, the clientless session grants access only to processes initiated by the web browser. For non-browser applications, an administrator can choose to tunnel all applications and thus remove the need to know which applications an end user may invoke. Table 91-6 shows in which situations processes are granted access.

Table 91-6 Access for Smart Tunnel Applications and Enabled Bookmarks

	Smart Tunnel Enabled Bookmark	Smart Tunnel Application Access
Application list specified	Any processes that match a process name in the application list are granted access.	Only processes that match a process name in the application list are granted access.
Smart tunnel is disabled	All processes (and their child processes) are granted access.	No process is granted access.
Smart Tunnel all Applications check box is checked	All processes (and their child processes) are granted access. **Note** This includes processes initiated by non-Smart Tunnel web pages if the web page is served by the same browser process.	All processes owned by the user who started the browser are granted access but not child processes of those original processes.

Restrictions

This configuration is applicable to Windows platforms only.

Detailed Steps

Follow these steps to configure tunnel policy.

Step 1 Choose **Configuration > Remote Access VPN > AAA/Local Users > Local Users**.

Step 2 In the User Account window, highlight the username that you want to edit.

Step 3 Click **Edit**. The Edit User Account window appears.

Step 4 In the left sidebar of the Edit User Account window, click **VPN Policy > Clientless SSL VPN**.

Step 5 Perform one of the following:

- Check the **smart tunnel_all_applications** check box. All applications will be tunneled without making a list or knowing which executables an end user may invoke for external applications.
- Or choose from the following tunnel policy options:
  - Uncheck the **Inherit** check box at the Smart Tunnel Policy parameter.

– Choose from the network list and specify one of the tunnel options: use smart tunnel for the specified network, do not use smart tunnel for the specified network, or use tunnel for all network traffic.

Adding Applications to Be Eligible for Smart Tunnel Access

The clientless SSL VPN configuration of each ASA supports *smart tunnel lists*, each of which identifies one or more applications eligible for smart tunnel access. Because each group policy or username supports only one smart tunnel list, you must group each set of applications to be supported into a smart tunnel list.

The Add or Edit Smart Tunnel entry dialog box lets you specify the attributes of an application in a smart tunnel list.

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels**, and choose a Smart Tunnel application list to edit, or add a new one.

Step 2 For a new list, enter a unique name for the list of applications or programs. 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.

Step 3 Click Add and add as many applications as you need to this smart tunnel list. The parameters are described below:

- Application ID - Enter a string to name the entry in the smart tunnel list. This user-specified name is saved and then returned onto the GUI. The string is unique for the operating system. 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 operating system, 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.

- OS - Click **Windows** or **Mac** to specify the host operating system of the application.
- Hash (Optional and only applicable to 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.

Step 4 Click **OK** to save the application, and create how ever many applications you need for this smart tunnel list.

Step 5 When you are done creating your 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 91-7 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 91-7 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

About Smart Tunnel Lists

For each group policy and username, you can configure clientless SSL VPN to do one of the following:

- Start smart tunnel access automatically upon user login.
- 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.

Restrictions

The smart tunnel logon options are mutually exclusive for each group policy and username. Use only one.

Creating a Smart Tunnel Auto Sign-On Server List

The Smart Tunnel Auto Sign-on Server List dialog box lets you add or edit lists of servers which will automate the submission of login credentials during smart tunnel setup. Auto Sign-on over a smart tunnel is available for Internet Explorer and Firefox.

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels**, and make sure Smart Tunnel Auto Sign-on Server List is expanded to display.

Step 2 Click **Add**, and enter a unique name for a list of remote servers that will help you to distinguish its contents or purpose from other lists that you are likely to configure. The string can be up to 64 characters. Do not use spaces.

After you create a smart tunnel auto sign-on list, that 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.

Adding Servers to a Smart Tunnel Auto Sign-on Server List

The following steps describe how to add servers to the list of servers for which to provide auto sign-on in smart tunnel connections, and assign that list to a group policies or a local user.

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels**, select one of the lists, and click **Edit**.

Step 2 Click the **Add** button on the Add Smart Tunnel Auto Sign-on Server List dialog to add one more more server smart tunnel servers.

Step 3 Enter the hostname or IP address of the server to auto-authenticate to:

- If you select Hostname, 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.
- If you select IP Address, enter an IP address.

Note Firefox does not support a host mask with wild cards, a subnet using IP addresses, or a netmask; you must use an exact host name or IP address. For example, within Firefox, if you enter *.cisco.com, auto sign-on to host email.cisco.com will fail.

Step 4 Windows Domain (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.

Step 5 HTTP-based Auto Sign-On (Optional)

- Authentication Realm - TheRealm is associated with the protected area of the website and passed back to the browser either in the authentication prompt or in the HTTP headers during authentication. Once auto-sign is configured here, and a realm string is specified, users can configure the realm string on a web application (such as Outlook Web Access) and access web applications without signing on.

  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 fails when the user clicks it.

Note If administrators do not know the corresponding realm, they should perform logon once and get the string from the prompt dialog.

- Port Number - Specify a port number for the corresponding hosts. For Firefox, if no port number is specified, auto sign-on is performed on HTTP and HTTPS, accessed by default port numbers 80 and 443 respectively.

Step 6 Click **OK**.

Step 7 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:
  1. Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies,** and open a group policy,
  2. Select the Portal tab, find the Smart Tunnel area, and choose the auto sign-on server list from the drop-down list next to the Auto Sign-on Server List attribute.
- To assign the list to a local user policy:
  1. choose **Configuration > Remote Access VPN > AAA/Local Users > Local Users**, and edit the local user you want to assign an auto sign-on server list to.
  2. Navigate to VPN Policy > Clientless SSL VPN, and find the Auto Sign-on Server setting under the Smart Tunnel area
  3. Uncheck Inherit, and choose a server list from the drop-down list next to the Auto Sign-on Server List attribute.

Enabling and Disabling Smart Tunnel Access

By default, smart tunnels are disabled.

If you enable smart tunnel access, the user will have to start it manually, using the **Application Access > Start Smart Tunnels** button on the clientless SSL VPN portal page.

Configuring Smart Tunnel Log Off

This section describes how to ensure that the smart tunnel is properly logged off. Smart tunnel can be logged off when all browser windows have been closed, or you can right click the notification icon and confirm log out.

Note We strongly recommend the use of the logout button on the portal. This method pertains to clientless SSL VPNs and logs off regardless of whether smart tunnel is used or not. The notification icon should be used only when using standalone applications without the browser.

When Its Parent Process Terminates

This practice requires the closing of all browsers to signify log off. The smart tunnel lifetime is now tied to the starting process lifetime. For example, if you started a smart tunnel from Internet Explorer, the smart tunnel is turned off when no iexplore.exe is running. Smart tunnel can determine that the VPN session has ended even if the user closed all browsers without logging out.

Note In some cases, a lingering browser process is unintentional and is strictly a result of an error. Also, when a Secure Desktop is used, the browser process can run in another desktop even if the user closed all browsers within the secure desktop. Therefore, smart tunnel declares all browser instances gone when no more visible windows exist in the current desktop.

With A Notification Icon

You may also choose to disable logging off when a parent process terminates so that a session survives if you close a browser. For this practice, you use a notification icon in the system tray to log out. The icon remains until the user clicks the icon to logout. If the session has expired before the user has logged out, the icon remains until the next connection is tried. You may have to wait for the session status to update in the system tray.

Note This icon is an alternative way to log out of SSL VPN. It is not an indicator of VPN session status.

Detailed Steps

To enable the icon in the notification area, follow these steps:

Step 1 Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels**.

Step 2 Enable the **Click on smart-tunnel logoff icon in the system tray** radio button.

Step 3 In the Smart Tunnel Networks portion of the window, check **Add** and enter both the IP address and hostname of the network which should include the icon.

Note If you right click the icon, a single menu item appears which prompts the user to log out of the SSL VPN.

Configuring Port Forwarding

The following sections describe port forwarding and how to configure it:

- Information About Port Forwarding, page 91-55
- Configuring DNS for Port Forwarding
- Adding Applications to Be Eligible for Port Forwarding
- Adding/Editing Port Forwarding Entry
- Assigning a Port Forwarding List
- Enabling and Disabling Port Forwarding

Information About Port Forwarding

Port forwarding lets users access TCP-based applications over a clientless SSL VPN connection. Such applications include the following:

- Lotus Notes
- Microsoft Outlook Express
- Perforce

- Sametime
- Secure FTP (FTP over SSH)
- SSH
- TELNET
- Windows Terminal Service

Other TCP-based applications may also work, but we have not tested them. Applications that use UDP do not work.

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.

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.

When configuring port forwarding on the ASA, you specify the port the application uses. When configuring smart tunnel access, you specify the name of the executable file or its path.

Prerequisites

- Refer to the Supported VPN Platforms, Cisco ASA 5500 Series compatibility guide for port forwarding pre-requisites.
- Browser-based users of Safari on Mac OS X 10.5.3 must identify a client certificate for use with the URL of the ASA, 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 or later 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 (or later) 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.

Restrictions

- 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 Microsoft Outlook Exchange (MAPI) proxy. However, you can configure smart tunnel support for Microsoft Office Outlook in conjunction with Microsoft Outlook Exchange Server.
- A stateful failover does not retain sessions established using Application Access (either port forwarding or smart tunnel access). Users must reconnect following a failover.

- 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.
- The port forwarding applet displays the local port and the remote port as the same when the local IP address 127.0.0.1 is being used and cannot be updated by the clientless SSL VPN connection from the ASA. As a result, the ASA creates new IP addresses 127.0.0.2, 127.0.0.3, and so on for local proxy IDs. Because you can modify the hosts file and use different loopbacks, the remote port is used as the local port in the applet. To connect, you can use Telnet with the host name, without specifying the port. The correct local IP addresses are available in the local hosts file.

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.

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 91-12).

Figure 91-12 Example DNS Server Values for Port Forwarding

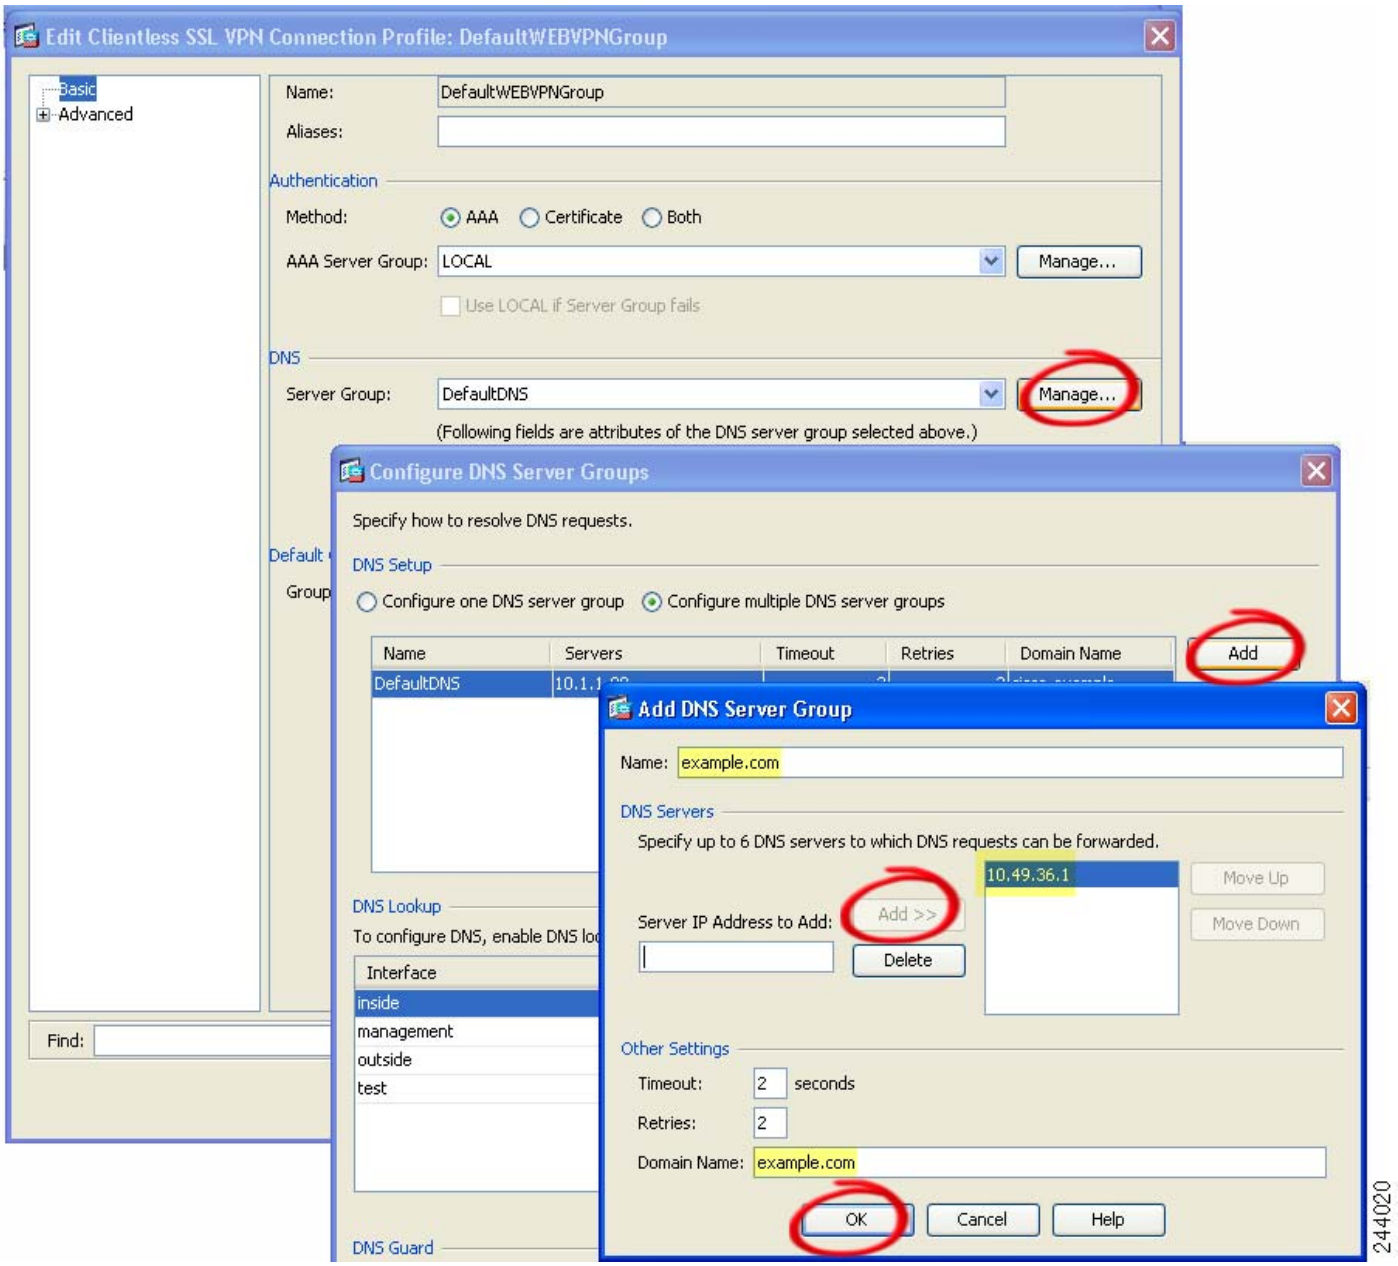

Note the domain name you entered. You need it when you specify the remote server later while configuring a port forwarding entry.

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**.

Adding Applications to Be Eligible for Port Forwarding

The clientless SSL VPN configuration of each ASA supports *port forwarding lists*, each of which specifies local and remote ports used by the applications for which you want to provide access. Because each group policy or username supports only one port forwarding list, you must group each set of applications to be supported into a list. To display the port forwarding list entries already present in the ASA configuration, enter the following commands:

Following the configuration of a port forwarding list, assign the list to group policies or usernames, as described in the next section.

Adding/Editing 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:

Prerequisites

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 the "Assigning a Port Forwarding List" section on page 91-59. The default setting for both the Domain and Server Group parameters is DefaultDNS.

Detailed Steps

Step 1 Click **Add**.

Step 2 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.

Step 3 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.

Step 4 Type the well-known port number for the application.

Step 5 Type a description of the application. The maximum is 64 characters.

Step 6 (Optional) Highlight a port forwarding list and click **Assign** to assign the selected list to one or more group policies, dynamic access policies, or user policies.

Assigning a Port Forwarding List

You can add or edit a named list of TCP applications to associate with users or group policies for access over clientless SSL VPN connections. For each group policy and username, you can configure clientless SSL VPN to do one of the following:

- Start port forwarding access automatically upon user login.
- Enable port forwarding access upon user login, but require the user to start it manually, using the **Application Access > Start Applications** button on the clientless SSL VPN Portal Page.

Note These options are mutually exclusive for each group policy and username. Use only one.

Detailed Steps

The Add or Edit Port Forwarding List dialog box lets you add or edit the following:

Step 1 Provide an alphanumeric name for the list. The maximum is 64 characters.

Step 2 Enter which local port listens for traffic for the application. 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.

Note Enter the IP address or DNS name 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.

Step 3 Enter the remote port that listens for traffic for the application.

Step 4 Describe the TCP application. The maximum is 64 characters.

For details, go to the section that addresses the option you want to use.

Enabling and Disabling Port Forwarding

By default, port forwarding is disabled.

If you enable port forwarding, the user will have to start it manually, using the **Application Access > Start Applications** button on the clientless SSL VPN portal page.

Configuring the Use of External Proxy Servers

Use the Proxies pane to configure the ASA 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.

Restrictions

HTTP and HTTPS proxy services do not support connections to personal digital assistants.

Detailed Steps

Step 1 Click Use an HTTP proxy server.

Step 2 Identify the HTTP proxy server by its IP address or hostname.

Step 3 Enter the hostname or IP address of the external HTTP proxy server.

Step 4 Enter the port that listens for HTTP requests. The default port is 80.

Step 5 (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.

Step 6 (Optional) Enter this keyword to accompany each HTTP proxy request with a username to provide basic, proxy authentication.

Step 7 Enter a password to send to the proxy server with each HTTP request.

Step 8 As an alternative to specifying the IP address of the HTTP proxy server, you can choose Specify PAC file URL 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 ASA ignores it.

Step 9 Choose if you want to use an HTTPS proxy server.

Step 10 Click to identify the HTTPS proxy server by its IP address or hostname.

Step 11 Enter the hostname or IP address of the external HTTPS proxy server.

Step 12 Enter the port that listens for HTTPS requests. The default port is 443.

Step 13 (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.

Step 14 (Optional) Enter this keyword to accompany each HTTPS proxy request with a username to provide basic, proxy authentication.

Step 15 Enter a password to send to the proxy server with each HTTPS request.

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.

Restrictions

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.

- Configuring SiteMinder and SAML Browser Post Profile, page 91-62—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 ASA 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 ASA 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 ASA, 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.

Detailed Steps

Use the SAML server documentation provided by the server software vendor to configure the SAML server in Relying Party mode. The following fields are displayed:

- 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 ASA currently supports the SiteMinder type and the SAML Browser Post Profile type.
- URL—*Display only.* Displays the SSO server URL to which the ASA 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 ASA 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.
- Assign—Highlight an SSO server and click this button to assign the selected server to one or more VPN group policies or user policies.

Step 1 Configure the SAML server parameters to represent the asserting party (the ASA):

- 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 ASA for SSO with SiteMinder, you must also configure your CA SiteMinder Policy Server with the Cisco authentication scheme, provided as a Java plug-in. 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:

Prerequisites

Configuring the SiteMinder Policy Server requires experience with SiteMinder.

Detailed 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/cisco/software/navigator.html and copy it to the default library directory for the SiteMinder server. This .jar file is also available on the Cisco ASA CD.

Adding or Editing 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.

Note NTLM is not the most secure choice and other better alternatives exist. NTLM does not support recent cryptographic methods such as AES or SHA-256. Cisco recommends that applications not use NTLM.

Detailed Steps

Step 1 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.

Step 2 *Display only.* Displays the type of SSO server. The types currently supported by the ASA are SiteMinder and SAML Browser Post Profile.

Step 3 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 ASA, the SSO server, and the SiteMinder Policy Server using the Cisco Java plug-in authentication scheme.

Step 4 Enter the number of times the ASA 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.

Step 5 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.

Application Access User Notes

The following sections provide information about using application access:

- Closing Application Access to Prevent hosts File Errors
- Closing Application Access to Prevent hosts File Errors
- Recovering from hosts File Errors When Using Application Access

Closing Application Access to Prevent hosts File Errors

To prevent hosts file errors that can interfere with Application Access, close the Application Access window properly when you finish using Application Access. To do so, click the close icon.

Recovering from hosts File Errors When Using Application Access

The following errors can occur if you do not close the Application Access window properly:

- The next time you try to start Application Access, it might be disabled; you receive a `Backup HOSTS File Found` error message.
- The applications themselves might be disabled or might malfunction, even when you are running them locally.

These errors can result from terminating the Application Access window in any improper way. For example:

- Your browser crashes while you are using Application Access.
- A power outage or system shutdown occurs while you are using Application Access.
- You minimize the Application Access window while you are working, then shut down your computer with the window active (but minimized).

This section includes the following topics:

- Understanding the hosts File
- Stopping Application Access Improperly
- Reconfiguring a Host's File Automatically Using Clientless SSL VPN
- Reconfiguring hosts File Manually

Understanding the hosts File

The hosts file on your local system maps IP addresses to host names. When you start Application Access, clientless SSL VPN modifies the hosts file, adding clientless SSL VPN-specific entries. Stopping Application Access by properly closing the Application Access window returns the file to its original state.

Before invoking Application Access...	hosts file is in original state.
When Application Access starts....	• Clientless SSL VPN copies the hosts file to hosts.webvpn, thus creating a backup. • Clientless SSL VPN then edits the hosts file, inserting clientless SSL VPN-specific information.
When Application Access stops...	• Clientless SSL VPN copies the backup file to the `hosts` file, thus restoring the hosts file to its original state. • Clientless SSL VPN deletes hosts.webvpn.
After finishing Application Access...	hosts file is in original state.

Note Microsoft anti-spyware software blocks changes that the port forwarding Java applet makes to the hosts file. See www.microsoft.com for information on how to allow hosts file changes when using anti-spyware software.

Stopping Application Access Improperly

When Application Access terminates abnormally, the `hosts` file remains in a clientless SSL VPN-customized state. Clientless SSL VPN checks the state the next time you start Application Access by searching for a hosts.webvpn file. If it finds one, a `Backup HOSTS File Found` error message (Figure 91-13) appears, and Application Access is temporarily disabled.

Once you shut down Application Access improperly, you leave your remote access client/server applications in limbo. If you try to start these applications without using clientless SSL VPN, they might malfunction. You might find that hosts that you normally connect to are unavailable. This situation could commonly occur if you run applications remotely from home, fail to quit the Application Access window before shutting down the computer, then try to run the applications later from the office.

Reconfiguring a Host's File Automatically Using Clientless SSL VPN

If you are able to connect to your remote access server, follow these steps to reconfigure the host's file and re-enable both Application Access and the applications.

Detailed Steps

Step 1 Start clientless SSL VPN and log in. The home page opens.

Step 2 Click the **Applications Access** then **Start Applications**, a `Backup HOSTS File Found` message appears. (See Figure 91-13.)

Figure 91-13 Backup HOSTS File Found Message

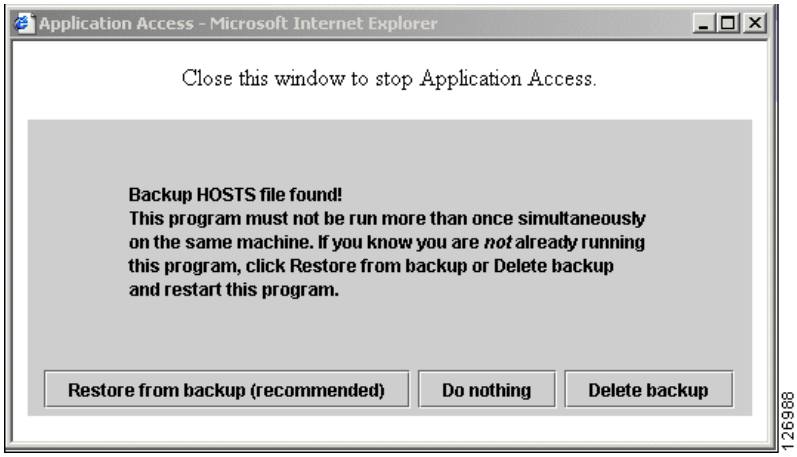

Step 3 Choose one of the following options:

- **Restore from backup**—Clientless SSL VPN forces a proper shutdown. It copies the hosts.webvpn backup file to the `hosts` file, restoring it to its original state, then deletes hosts.webvpn. You then have to restart Application Access.
- **Do nothing**—Application Access does not start. The remote access home page reappears.
- **Delete backup**—Clientless SSL VPN deletes the hosts.webvpn file, leaving the hosts file in its clientless SSL VPN-customized state. The original `hosts` file settings are lost. Application Access then starts, using the clientless SSL VPN-customized hosts file as the new original. Choose this

option only if you are unconcerned about losing hosts file settings. If you or a program you use might have edited the hosts file after Application Access has shut down improperly, choose one of the other options, or edit the hosts file manually. (See "Reconfiguring hosts File Manually.")

Reconfiguring hosts File Manually

If you are not able to connect to your remote access server from your current location, or if you have customized the hosts file and do not want to lose your edits, follow these steps to reconfigure the hosts file and reenable both Application Access and the applications.

Detailed Steps

Step 1 Locate and edit your hosts file. The most common location is c:\windows\sysem32\drivers\etc\hosts.

Step 2 Check to see if any lines contain the string: `# added by WebVpnPortForward`
If any lines contain this string, your hosts file is clientless SSL VPN-customized. If your hosts file is clientless SSL VPN-customized, it looks similar to the following example:

```
server1 # added by WebVpnPortForward
server1.example.com invalid.cisco.com # added by WebVpnPortForward
server2 # added by WebVpnPortForward
server2.example.com invalid.cisco.com # added by WebVpnPortForward
server3 # added by WebVpnPortForward
server3.example.com invalid.cisco.com # added by WebVpnPortForward

# Copyright (c) 1993-1999 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97     cisco.example.com          # source server
#       38.25.63.10     x.example.com              # x client host

127.0.0.1       localhost
```

Step 3 Delete the lines that contain the string: `# added by WebVpnPortForward`

Step 4 Save and close the file.

Step 5 Start clientless SSL VPN and log in.

The home page appears.

Step 6 Click the **Application Access** link.

The Application Access window appears. Application Access is now enabled.

Configuring File Access

Clientless SSL VPN serves remote users with HTTPS portal pages that interface with proxy CIFS and/or FTP clients running on the ASA. Using either CIFS or FTP, clientless SSL VPN provides users with network access to the files on the network, to the extent that the users meet user authentication requirements and the file properties do not restrict access. The CIFS and FTP clients are transparent; the portal pages delivered by clientless SSL VPN provide the appearance of direct access to the file systems.

When a user requests a list of files, clientless SSL VPN queries the server designated as the master browser for the IP address of the server containing the list. The ASA gets the list and delivers it to the remote user on a portal page.

Clientless SSL VPN lets the user invoke the following CIFS and FTP functions, depending on user authentication requirements and file properties:

- Navigate and list domains and workgroups, servers within a domain or workgroup, shares within a server, and files within a share or directory
- Create directories
- Download, upload, rename, move, and delete files

The ASA uses a master browser, WINS server, or DNS server, typically on the same network as the ASA or reachable from that network, to query the network for a list of servers when the remote user clicks **Browse Networks** in the menu of the portal page or on the toolbar displayed during the clientless SSL VPN session.

The master browser or DNS server provides the CIFS/FTP client on the ASA with a list of the resources on the network, which clientless SSL VPN serves to the remote user.

Note Before configuring file access, you must configure the shares on the servers for user access.

CIFS File Access Requirement and Limitation

To access `\\server\share\subfolder\personal folder`, the user must have a minimum of read permission for all parent folders, including the share itself.

Use **Download** or **Upload** to copy and paste files to and from CIFS directories and the local desktop. The Copy and Paste buttons are intended for remote to remote actions only, not local to remote, or remote to local.

The CIFS browse server feature does not support double-byte character share names (share names exceeding 13 characters in length). This only affects the list of folders displayed, and does not affect user access to the folder. As a workaround, you can pre-configure the bookmark(s) for the CIFS folder(s) that use double-byte share names, or the user can enter the URL or bookmark of the folder in the format cifs://server/<long-folder-name> . For example:

```
cifs://server/Do you remember?
cifs://server/Do%20you%20remember%3F
```

Adding Support for File Access

Configure file access as follows:

Note The first procedure describes how to specify the master browser and WINS servers. As an alternative, you can use ASDM to configure URL lists and entries that provide access to file shares.

Adding a share in ASDM does not require a master browser or a WINS server. However, it does not provide support for the Browse Networks link. You can use a hostname or an IP address to refer to ServerA when entering this command. If you use a hostname, the ASA requires a DNS server to resolve it to an IP address.

For a complete description of these commands, see the *Cisco Security Appliance Command Reference*.

Ensuring Clock Accuracy for SharePoint Access

The clientless SSL VPN server on the ASA uses cookies to interact with applications such as Microsoft Word on the endpoint. The cookie expiration time set by the ASA can cause Word to malfunction when accessing documents on a SharePoint server if the time on the ASA is incorrect. To prevent this malfunction, set the ASA clock properly. We recommend configuring the ASA to dynamically synchronize the time with an NTP server. For instructions, see "Setting the Date and Time."

Customizing the Clientless SSL VPN User Experience

You can customize the clientless SSL VPN user experience, including the logon, portal, and logout pages. There are two methods you can use. You can customize pre-defined page components in the Add/Edit Customization Object window. This window adds, or makes changes to, an XML file stored on the ASA (a customization object) that is used to customize the pages. Alternatively, you can export the XML file to a local computer or server, make changes to the XML tags, and re-import the file to the ASA. Either method creates a customization object that you apply to a connection profile or group policy.

Rather than customizing the pre-defined components of the logon page, you can create your own page and import it to the ASA for full customization. To do this see the "Replacing the Logon Page with your own Fully Customized Page" section on page 91-71.

You can customize pre-defined components of the logon page, including titles, language options, and messages to users. Alternatively, you can completely replace the page with your own custom page (full customization). The following sections detail both procedures:

- Customizing the Logon Page with the Customization Editor, page 91-69
- Replacing the Logon Page with your own Fully Customized Page, page 91-71

Customizing the Logon Page with the Customization Editor

Figure 91-14 shows the logon page and the pre-defined components you can customize:

Figure 91-14 Components of Clientless Logon Page

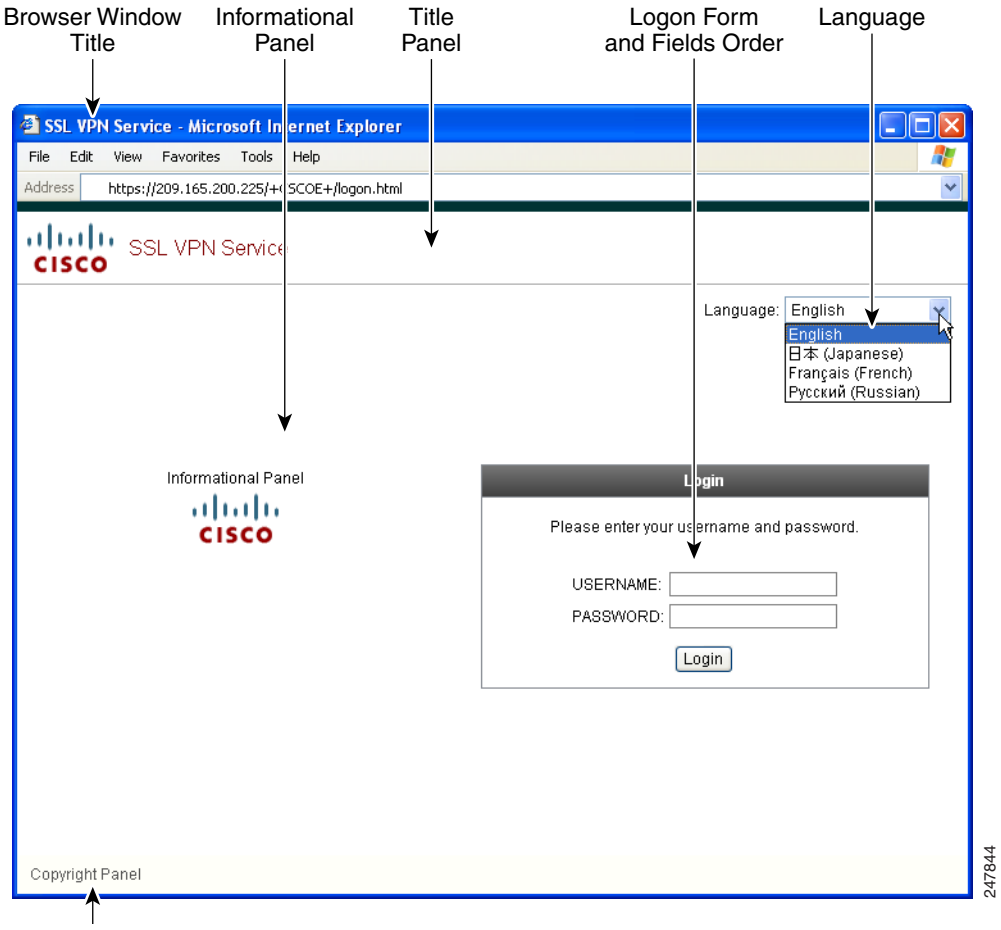

To customize all the components of the logon page, follow this procedure. You can preview your changes for each component by clicking the Preview button:

Step 1 Specify pre-defined customization. Go to Logon Page and select **Customize pre-defined logon page components**. Specify a title for the browser window.

Step 2 Display and customize the title panel. Go to Logon Page > Title Panel and check **Display title panel**. Enter text to display as the title and specify a logo. Specify any font styles.

Step 3 Specify language options to display. Go to Logon Page > Language and check **Enable Language Selector**. Add or delete any languages to display to remote users. Languages in the list require translation tables that you configure in Configuration > Remote Access VPN > Language Localization.

Step 4 Customize the logon form. Go to Logon Page > Logon Form. Customize the text of the form and the font style in the panel. The secondary password field appears to users only if a secondary authentication server is configured in the connection profile.

Step 5 Arrange the position of the logon form fields. Go to Logon Page > Form Fields Order. Use the up and down arrow buttons to change the order that the fields are displayed.

Step 6 Add messages to users. Go to Logon Page > Informational Panel and check **Display informational panel**. Add text to display in the panel, change the position of the panel relative to the logon form, and specify a logo to display in this panel.

Step 7 Display a copyright statement. Go to Logon Page > Copyright Panel and check **Display copyright panel**. Add text to display for copyright purposes.

Step 8 Click **OK**, then apply the changes to the customization object you edited.

Replacing the Logon Page with your own Fully Customized Page

If you prefer to use your own, custom login screen, rather than changing specific components of the logon page we provide, you can perform this advanced customization using the Full Customization feature.

With Full Customization, you provide the HTML for your own login screen, and you insert Cisco HTML code that calls functions on the ASA that create the Login form and the Language Selector drop-down list.

This document describes the modifications you need to make to your HTML code and the tasks required to configure the ASA to use your code.

Figure 91-15 shows a simple example of a custom login screen enabled by the Full Customization feature.

Figure 91-15 Example of Full Customization of Logon Page

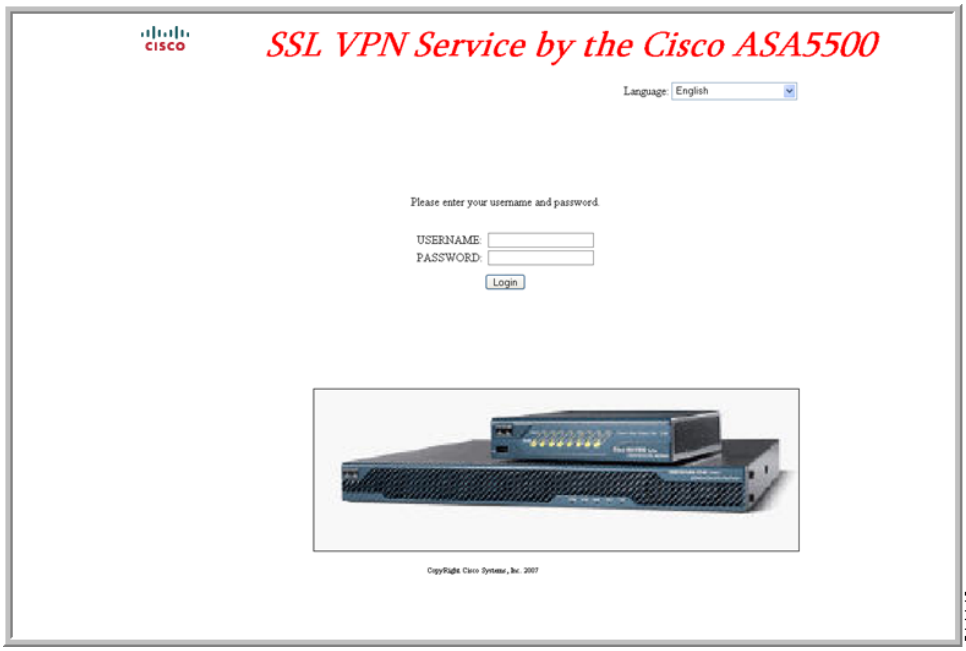

The following sections describe the tasks to customize the login screen:

- Creating the Custom Login Screen File
- Importing the File and Images
- Configuring the Security Appliance to use the Custom Login Screen

Creating the Custom Login Screen File

The following HTML code is used as an example and is the code that displays:

```
<head>
<meta http-equiv="Content-Type" content="text/html; charset=windows-1252">
<title>New Page 3</title>
<base target="_self">
</head>

<p align="center">
<img border="0" src="/+CSCOU+/cisco_logo.jpg" width="188" height="48"><font face="Snap
ITC" size="6" color="#FF00FF">
</font><font face="Snap ITC" color="#FF00FF" size="7"> </font><i><b><font
color="#FF0000" size="7" face="Sylfaen"> SSL VPN Service by the Cisco
ASA5500</font></b></i></p>

    <body onload="csco_ShowLoginForm('lform');csco_ShowLanguageSelector('selector')">

    <table>

    <tr><td colspan=3 height=20 align=right><div id="selector" style="width:
    300px"></div></td></tr>
    <tr><td></td><td></td><td></td></tr>
    <tr>
    <td height="379"></td>
    <td height="379"></td>
    <td align=middle valign=middle>
    <div id=lform >
    <p> </p>
    <p> </p>
    <p> </p>
    <p>Loading...</p>
    </div>
    </td>
    </tr>
    <tr>
    <td width="251"></td>
    <td width="1"></td>
    <td align=right valign=right width="800">
    <img border="1" src="/+CSCOU+/asa5500.jpg" width="660" height="220" align="middle">
    </td></tr>

    </table>
```

The indented code injects the Login form and the Language Selector on the screen. The function **csco_ShowLoginForm('lform')** injects the logon form. **csco_ShowLanguageSelector('selector')** injects the Language Selector.

Follow these steps to modify your HTML file:

Step 1 Name your file **logon.inc**. When you import the file, the ASA recognizes this filename as the logon screen.

Step 2 Modify the paths of images used by the file to include **/+CSCOU+/**.

Files that are displayed to remote users before authentication must reside in a specific area of the ASA cache memory represented by the path /+CSCOU+/. Therefore, the source for each image in the file must include this path. For example:

```
src="/+CSCOU+/asa5520.gif"
```

Step 3 Insert the special HTML code below. This code contains the Cisco functions, described earlier, that inject the login form and language selector onto the screen.

```
<body onload="csco_ShowLoginForm('lform');csco_ShowLanguageSelector('selector')">

<table>

<tr><td colspan=3 height=20 align=right><div id="selector" style="width:
300px"></div></td></tr>
<tr><td></td><td></td><td></td></tr>
<tr>
<td height="379"></td>
<td height="379"></td>
<td align=middle valign=middle>
<div id=lform >
<p> </p>
<p> </p>
<p> </p>
<p>Loading...</p>
</div>
</td>
</tr>
<tr>
<td width="251"></td>
<td width="1"></td>
<td align=right valign=right width="800">
<img border="1" src="/+CSCOU+/asa5500.jpg" width="660" height="220" align="middle">
</td></tr>

</table>
```

Importing the File and Images

Follow these steps to import your HTML file and any images to the ASA:

Step 1 Go to **Clientless SSL VPN Access > Portal > Web Contents**.

Step 2 Click Import. The Import Web Content window displays.

- **a.** Select the **Source** option, and enter the path the web content files.
- **b.** In the **Destination** area, select **No** for *Require Authentication to access its content* . This ensures the files are stored in the area of flash memory accessible to users before authentication.

Step 3 Click Import Now, and kee used by the file as Web Content using the same window.

Configuring the Security Appliance to use the Custom Login Screen

Follow these steps to enable the ASA to use the new login screen in a customization object:

Step 1 Select a customization object. Go to **Clientless SSL VPN Access > Portal > Customization**. Select a customization object in the table and click **Edit**. The Edit Customization Object window displays.

Step 2 In the navigation pane, select Logon Page.

Step 3 Choose Replace pre-defined logon page with a custom page.

Step 4 Click Manage to import your logon page file. The Import Web Content window displays.

Step 5 In the Destination area, select **No** to ensure your logon page is visible to users before they authenticate.

Step 6 Back in the Edit Customization Object window, click General and enable the customization object for the connection profile and/or group policies you desire.

Using Clientless SSL VPN with PDAs

You can access clientless SSL VPN from your Pocket PC or other certified personal digital assistant device. Neither the ASA administrator nor the clientless SSL VPN user need do anything special to use clientless SSL VPN with a certified PDA.

Cisco has certified the following PDA platform:

HP iPaq H4150
Pocket PC 2003
Windows CE 4.20.0, build 14053
Pocket Internet Explorer (PIE)
ROM version 1.10.03ENG
ROM Date: 7/16/2004

Some differences in the PDA version of clientless SSL VPN exist:

- A banner web page replaces the popup clientless SSL VPN window.
- An icon bar replaces the standard clientless SSL VPN floating toolbar. This bar displays the Go, Home and Logout buttons.
- The Show Toolbar icon is not included on the main clientless SSL VPN portal page.
- Upon clientless SSL VPN logout, a warning message provides instructions for closing the PIE browser properly. If you do not follow these instructions and you close the browser window in the common way, PIE does not disconnect from clientless SSL VPN or any secure website that uses HTTPS.

Restrictions

- Clientless SSL VPN supports OWA 2000 and OWA 2003 Basic Authentication. If Basic Authentication is not configured on an OWA server and a clientless SSL VPN user attempts to access that server, access is denied.
- Unsupported clientless SSL VPN features:
  - Application Access and other Java-dependent features.
  - HTTP proxy.
  - The Citrix Metaframe feature (if the PDA does not have the corresponding Citrix ICA client software).

Using E-Mail over Clientless SSL VPN

Clientless SSL VPN supports several ways to access e-mail. This section includes the following methods:

- Configuring E-mail Proxies

- Configuring Web E-mail: MS Outlook Web App

Configuring E-mail Proxies

Clientless SSL VPN supports IMAP4S, POP3S, and SMTPS e-mail proxies. The following attributes apply globally to e-mail proxy users.

Restrictions

E-mail clients such as MS Outlook, MS Outlook Express, and Eudora lack the ability to access the certificate store.

1 With the Eudora e-mail client, SMTPS works only on port 465, even though the default port for SMTPS connections is 988.

Configuring Web E-mail: MS Outlook Web App

The ASAsupports Microsoft Outlook Web App to Exchange Server 2010 and Microsoft Outlook Web Access to Exchange Server 2007, 2003, and 2000. OWA requires that users perform the following steps:

Detailed Steps

Step 1 Enter the URL of the e-mail service into the address field or click an associated bookmark in the clientless SSL VPN session.

Step 2 When prompted, enter the e-mail server username in the format *domain\username*.

Step 3 Enter the e-mail password.

Configuring Portal Access Rules

This enhancement allows customers to configure a global clientless SSL VPN access policy to permit or deny clientless SSL VPN sessions based on the data present in the HTTP header. If the ASA denies a clientless SSL VPN session, it returns an error code to the endpoint immediately.

The ASA evaluates this access policy before the endpoint authenticates to the ASA. As a result, in the case of a denial, fewer ASA processing resources are consumed by additional connection attempts from the endpoint.

Prerequisites

Log on to the ASA and enter global configuration mode. In global configuration mode, the ASA displays this prompt:

```
hostname(config)#
```

Detailed Steps

Step 1 Start ASDM and select **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Portal Access Rule**.

The Portal Access Rule window opens.

Step 2 Click **Add** to create a portal access rule or select an existing rule and click **Edit.**

The Add (or Edit) Portal Access Rule dialog box opens.

Step 3 Enter a rule number from 1-65535 in the Rule Priority field.

Rules are processed in order of priority from 1-65535.

Step 4 In the User Agent field, enter the name of the user agent you want to find in the HTTP header.

- Surround the string with wildcards (*) to generalize the string; for example, *Thunderbird*. We recommend using wildcards in your search string. Without wildcards, the rule may not match any strings or it may match many fewer strings than you expect.
- If your string contains a space, ASDM automatically adds quotes to the beginning and end of the string when it saves the rule. For example, if you enter `my agent`, ASDM will save the string as `"my agent"`. ASA will then search for matches of `my agent`.

  Do not add quotes to a string with spaces yourself unless you want ASA to match the quotes you added to the string. For example, if you enter `"my agent"` ASDM will save the string as `"\"my agent\""` and try to find a match for `"my agent"` and it will not find `my agent`.
- If you want to use wildcards with a string that contains a space, start and end the entire string with wildcards, for example, `*my agent*` and ASDM will automatically surround that string with quotes when it saves the rule.

Step 5 In the Action field, select either **Deny** or **Permit**.

The ASA will deny or permit a clientless SSL VPN connection based on this setting.

Step 6 Enter an HTTP message code in the Returned HTTP Code field.

The HTTP message number 403 is pre-populated in the field and is the default value for portal access rules. The allowed range of message codes is 200-599.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Using Proxy Bypass

You can configure the ASA 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 ASA. 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 everything in a URL after the .com or .org or other types of 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*.

Detailed Steps

You can set rules for when the ASA performs little or no content rewriting:

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Proxy Bypass**.

Step 2 Select the Interface name for proxy bypass.

Step 3 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 ASA 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.

Step 4 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.

Step 5 Specify 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.

Clientless SSL VPN End User Setup

This section is for the system administrator who sets up clientless SSL VPN for end users. It describes how to customize the end-user interface.

This section summarizes configuration requirements and tasks for a remote system. It specifies information to communicate to users to get them started using clientless SSL VPN. It includes the following topics:

- Defining the End User Interface
- Customizing Clientless SSL VPN Pages, page 91-80
- Customizing Help, page 91-121
- Requiring Usernames and Passwords
- Communicating Security Tips
- Configuring Remote Systems to Use Clientless SSL VPN Features

Defining the End User Interface

The clientless SSL VPN end user interface consists of a series of HTML panels. A user logs on to clientless SSL VPN by entering the IP address of an ASA interface in the format https://*address*. The first panel that displays is the login screen (Figure 91-16).

Figure 91-16 Clientless SSL VPN Login Screen

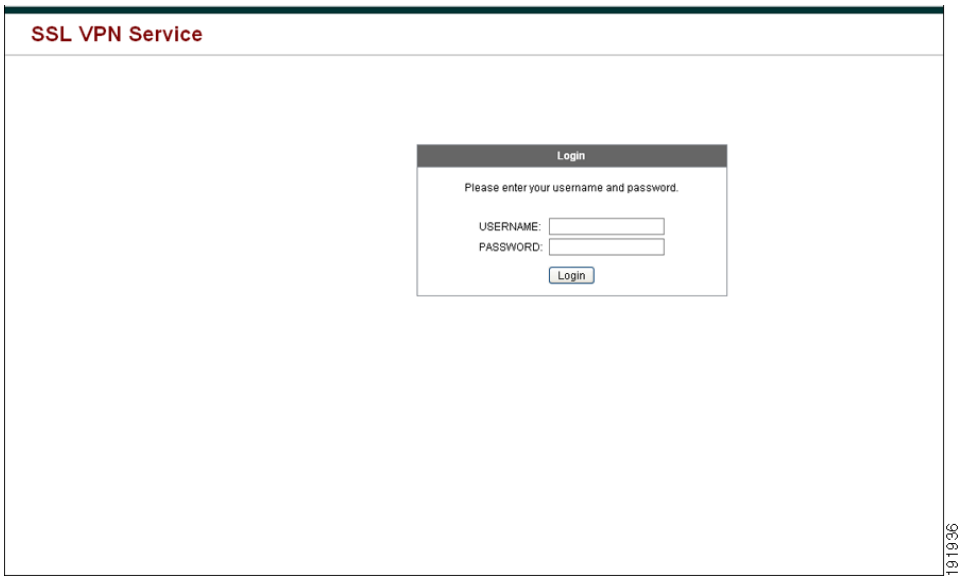

Viewing the Clientless SSL VPN Home Page

After the user logs in, the portal page opens.

The home page displays all of the clientless SSL VPN features you have configured, and its appearance reflects the logo, text, and colors you have selected. This sample home page includes all available clientless SSL VPN features with the exception of identifying specific file shares. It lets users browse the network, enter URLs, access specific websites, and use Application Access (port forwarding and smart tunnels) to access TCP applications.

Viewing the Clientless SSL VPN Application Access Panel

To start port forwarding or smart tunnels, a user clicks the **Go** button in the Application Access box. The Application Access window opens (Figure 91-17).

Figure 91-17 Clientless SSL VPN Application Access Window

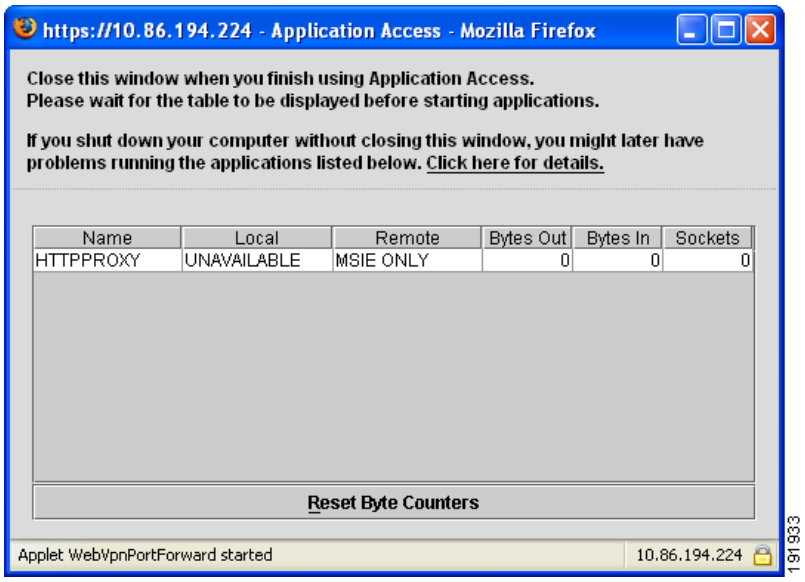

This window displays the TCP applications configured for this clientless SSL VPN connection. To use an application with this panel open, the user starts the application in the normal way.

Note A stateful failover does not retain sessions established using Application Access. Users must reconnect following a failover.

Viewing the Floating Toolbar

The floating toolbar shown in Figure 91-18 represents the current clientless SSL VPN session.

Figure 91-18 Clientless SSL VPN Floating Toolbar

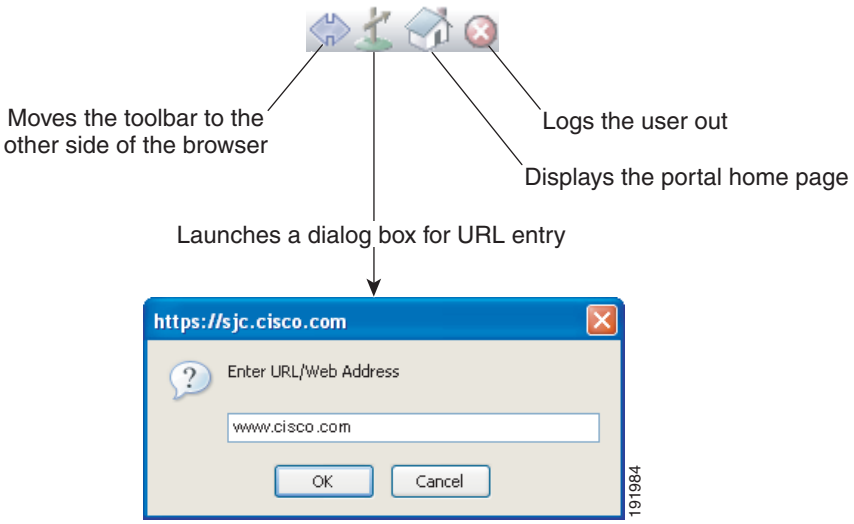

Be aware of the following characteristics of the floating toolbar:

- 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.
- If you close the toolbar, the ASA prompts you to confirm that you want to end the clientless SSL VPN session.

See Table 91-14 on page 91-123 for detailed information about using clientless SSL VPN.

Customizing Clientless SSL VPN Pages

You can change the appearance of the portal pages displayed to clientless SSL VPN users. This includes the Login page displayed to users when they connect to the security appliance, the Home page displayed to users after the security appliance authenticates them, the Application Access window displayed when users launch an application, and the Logout page displayed when users log out of clientless SSL VPN sessions.

After you customize the portal pages, you can save your customization and apply it to a specific connection profile, group policy, or user. The changes do not take effect until you reload the ASA, or you disable and then enable clientless SSL.

You can create and save many customization objects, enabling the security appliance to change the appearance of portal pages for individual users or groups of users.

This section includes the following topics:

- Information About Customization, page 91-80
- Exporting a Customization Template, page 77-84
- Editing the Customization Template, page 91-81
- Login Screen Advanced Customization, page 91-87
- Login Screen Advanced Customization, page 91-87

Information About Customization

The ASA uses customization objects to define the appearance of user screens. A customization object is compiled from an XML file which contains XML tags for all the customizable screen items displayed to remote users. The ASA software contains a customization template that you can export to a remote PC. You can edit this template and import the template back into the ASA as a new customization object.

When you export a customization object, an XML file containing XML tags is created at the URL you specify. The XML file created by the customization object named *Template* contains empty XML tags, and provides the basis for creating new customization objects. This object cannot be changed or deleted from cache memory but can be exported, edited, and imported back into the ASA as a new customization object.

Customization Objects, Connection Profiles, and Group Policies

Initially, when a user first connects, the default customization object (named *DfltCustomization*) identified in the connection profile (tunnel group) determines how the logon screen appears. If the connection profile list is enabled, and the user selects a different group which has its own customization, the screen changes to reflect the customization object for that new group.

After the remote user is authenticated, the screen appearance is determined by whether a customization object that has been assigned to the group policy.

Exporting a Customization Template

When you export a customization object, an XML file is created at the URL you specify. The customization template (named *Template*) contains empty XML tags and provides the basis for creating new customization objects. This object cannot be changed or deleted from cache memory but can be exported, edited, and imported back into the ASA as a new customization object.

Editing the Customization Template

This section shows the contents of the customization template and has convenient figures to help you quickly choose the correct XML tag and make changes that affect the screens.

You can use a text editor or an XML editor to edit the XML file. The following example shows the XML tags of the customization template. Some redundant tags have been removed for easier viewing:

Example:

```
<custom>
   <localization>
      <languages>en,ja,zh,ru,ua</languages>
      <default-language>en</default-language>
   </localization>
    <auth-page>
      <window>
          <title-text l10n="yes"><![CDATA[SSL VPN Service]]></title-text>
      </window>
      <full-customization>
          <mode>disable</mode>
          <url></url>
      </full-customization>
      <language-selector>
        <mode>disable</mode>
        <title l10n="yes">Language:</title>
        <language>
          <code>en</code>
          <text>English</text>
        </language>
        <language>
          <code>zh</code>
          <text>ä¸-å›½ (Chinese)</text>
        </language>
        <language>
          <code>ja</code>
          <text>æ—¥æœ¬ (Japanese)</text>
        </language>
        <language>
          <code>ru</code>
          <text>Ð ÑƒÑÑкий (Russian)</text>
        </language>
        <language>
          <code>ua</code>
          <text>Ð£ÐºÑ Ð°Ñ—Ð½ÑÑŒÐºÐ° (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:</internal-password-prompt-text>
            <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>
            <border-color>#858A91</border-color>
        </logon-form>
        <logout-form>
            <title-text l10n="yes"><![CDATA[Logout]]></title-text>
            <message-text l10n="yes"><![CDATA[Goodbye.<br>

For your own security, please:<br>

<li>Clear the browser's cache

<li>Delete any downloaded files

<li>Close the browser's window]]></message-text>
            <login-button-text l10n="yes">Logon</login-button-text>
            <hide-login-button>no</hide-login-button>
            <title-background-color><![CDATA[#666666]]></title-background-color>
            <title-font-color><![CDATA[#ffffff]]></title-font-color>
            <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>
            <border-color>#858A91</border-color>
        </logout-form>
        <title-panel>
           <mode>enable</mode>
           <text l10n="yes"><![CDATA[SSL VPN Service]]></text>
           <logo-url l10n="yes">/+CSCOU+/csco_logo.gif</logo-url>
           <gradient>yes</gradient>
           <style></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"></text>
        </info-panel>
        <copyright-panel>
           <mode>disable</mode>
           <text l10n="yes"></text>
        </copyright-panel>
    </auth-page>
    <portal>
      <title-panel>
         <mode>enable</mode>
```

```
   <text l10n="yes"><![CDATA[SSL VPN Service]]></text>
   <logo-url l10n="yes">/+CSCOU+/csco_logo.gif</logo-url>
   <gradient>yes</gradient>
   <style></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></title>
   <text></text>
```

```
            <notitle></notitle>
            <column></column>
            <row></row>
            <height></height>
         </pane>
         <pane>
            <type>IMAGE</type>
            <mode>disable</mode>
            <title></title>
            <url l10n="yes"></url>
            <notitle></notitle>
            <column></column>
            <row></row>
            <height></height>
         </pane>
         <pane>
            <type>HTML</type>
            <mode>disable</mode>
            <title></title>
            <url l10n="yes"></url>
            <notitle></notitle>
            <column></column>
            <row></row>
            <height></height>
         </pane>
         <pane>
            <type>RSS</type>
            <mode>disable</mode>
            <title></title>
            <url l10n="yes"></url>
            <notitle></notitle>
            <column></column>
            <row></row>
            <height></height>
         </pane>
         <url-lists>
            <mode>group</mode>
         </url-lists>
         <home-page>
            <mode>standard</mode>
            <url></url>
         </home-page>
      </portal>
</custom>
```

Figure 91-19 shows the Logon page and its customizing XML tags. All these tags are nested within the higher-level tag <auth-page>.

Figure 91-19 Logon Page and Associated XML Tags

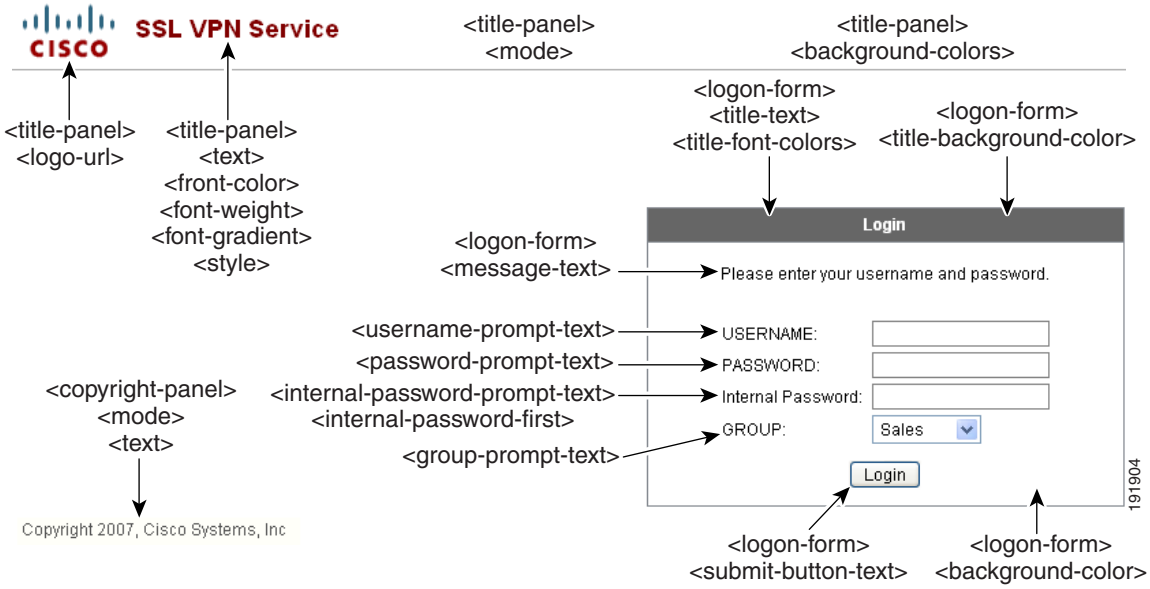

Figure 91-20 shows the Language Selector drop-down list that is available on the Logon page, and the XML tags for customizing this feature. All these tags are nested within the higher-level <auth-page> tag.

Figure 91-20 Language Selector on Logon Screen and Associated XML Tags

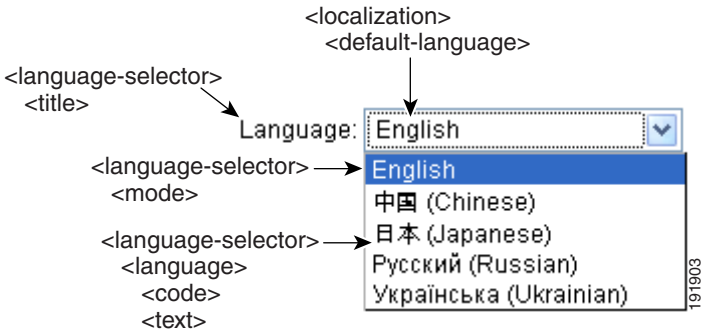

Figure 91-21 shows the Information Panel that is available on the Logon page, and the XML tags for customizing this feature. This information can appear to the left or right of the login box. These tags are nested within the higher-level <auth-page> tag.

Figure 91-21 Information Panel on Logon Screen and Associated XML Tags

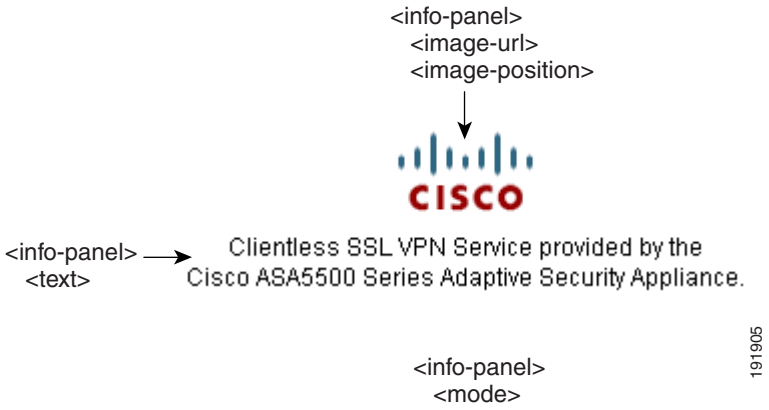

Figure 91-22 shows the Portal page and the XML tags for customizing this feature. These tags are nested within the higher-level <auth-page> tag.

Figure 91-22 Portal Page and Associated XML Tags

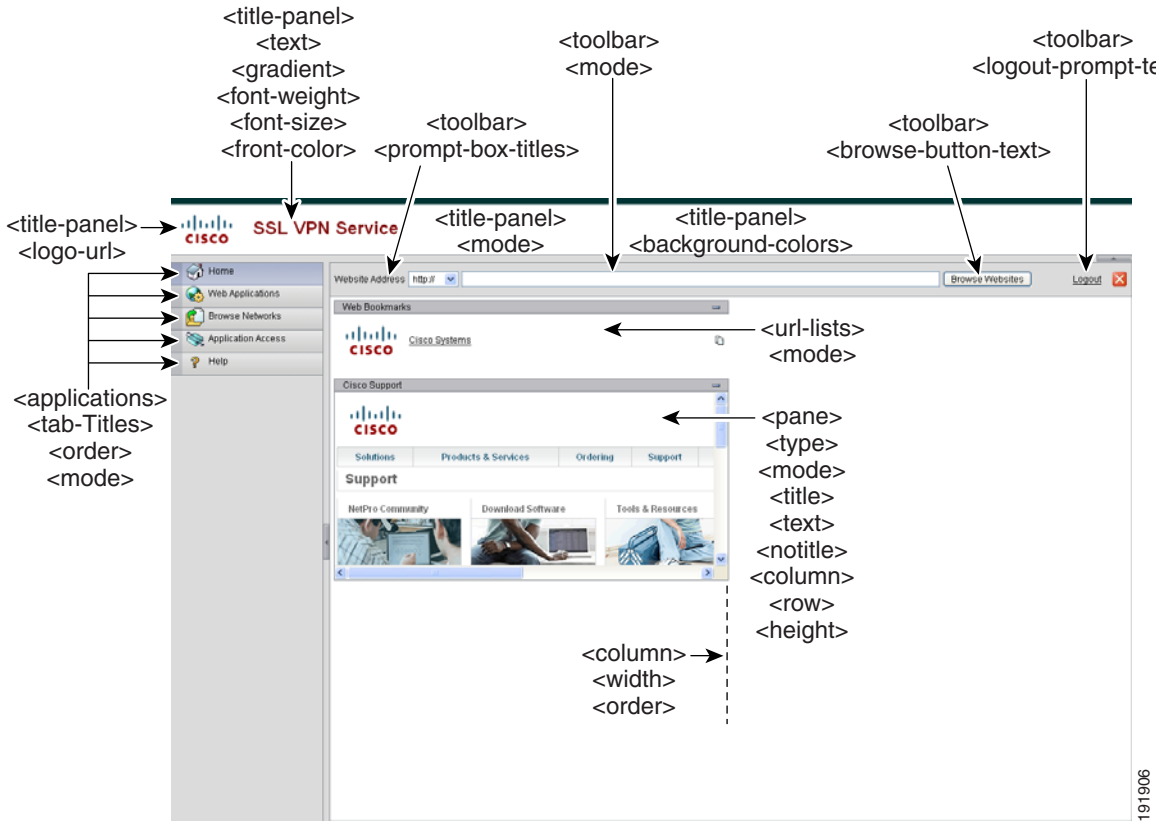

Login Screen Advanced Customization

If you prefer to use your own, custom login screen, rather than changing specific screen elements of the login screen we provide, you can perform this advanced customization using the *Full Customization* feature.

With Full Customization, you provide the HTML for your own login screen, and you insert Cisco HTML code that calls functions on the ASA that create the Login form and the Language Selector drop-down list.

This section describes the modifications you need to make to your HTML code and the tasks required to configure the ASA to use your code.

Figure 91-23 shows the standard Cisco login screen that displays to clientless SSL VPN users. The Login form is displayed by a function called by the HTML code.

Figure 91-23 Standard Cisco Login Page

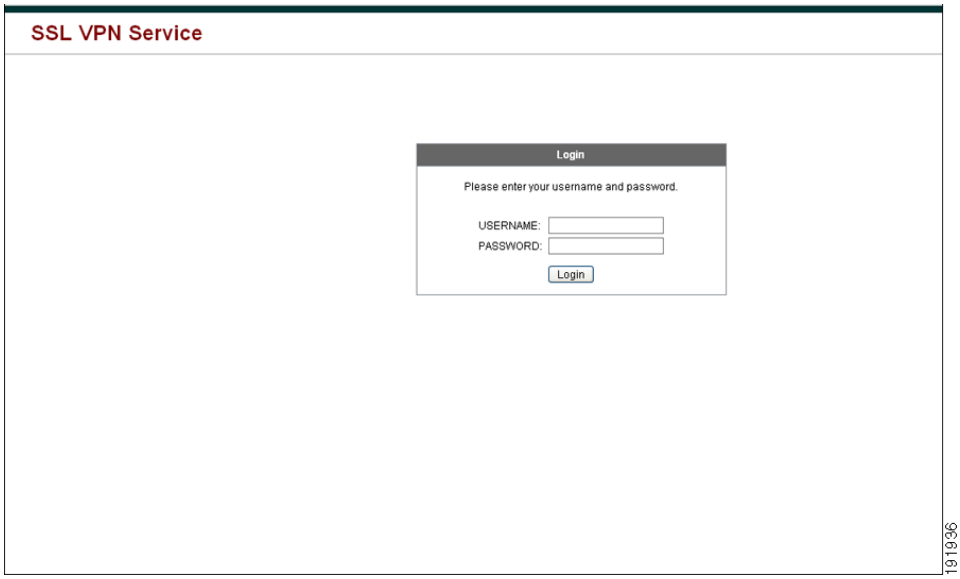

Figure 91-24 shows the Language Selector drop-down list. This feature is an option for clientless SSL VPN users and is also called by a function in the HTML code of the login screen.

Figure 91-24 Language Selector Drop-down List

Figure 91-25 shows a simple example of a custom login screen enabled by the Full Customization feature.

Figure 91-25 Example of Full Customization of Login Screens

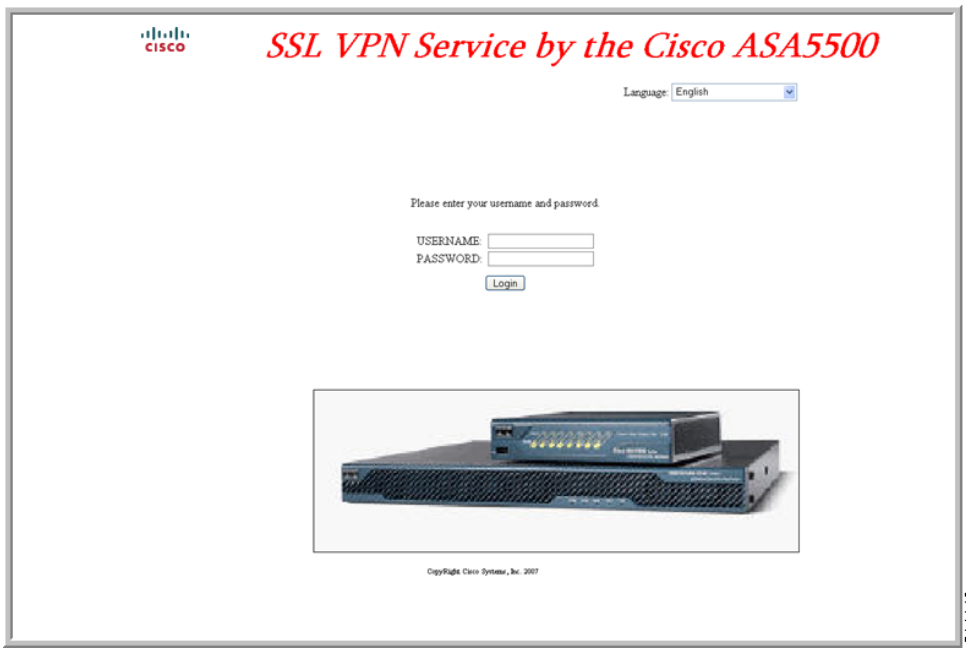

The following HTML code is used as an example and is the code that displays:

Example:

```
<head>
<meta http-equiv="Content-Type" content="text/html; charset=windows-1252">
<title>New Page 3</title>
<base target="_self">
</head>

<p align="center">
<img border="0" src="/+CSCOU+/cisco_logo.jpg" width="188" height="48"><font face="Snap
ITC" size="6" color="#FF00FF">
</font><font face="Snap ITC" color="#FF00FF" size="7"> </font><i><b><font
color="#FF0000" size="7" face="Sylfaen"> SSL VPN Service by the Cisco
ASA5500</font></b></i></p>

    <body onload="csco_ShowLoginForm('lform');csco_ShowLanguageSelector('selector')">

    <table>

    <tr><td colspan=3 height=20 align=right><div id="selector" style="width:
    300px"></div></td></tr>
    <tr><td></td><td></td><td></td></tr>
    <tr>
    <td height="379"></td>
    <td height="379"></td>
    <td align=middle valign=middle>
    <div id=lform >
    <p> </p>
```

```
<p> </p>
<p> </p>
<p>Loading...</p>
</div>
</td>
</tr>
<tr>
<td width="251"></td>
<td width="1"></td>
<td align=right valign=right width="800">
<img border="1" src="/+CSCOU+/asa5500.jpg" width="660" height="220" align="middle">
</td></tr>

</table>
```

The indented code injects the Login form and the Language Selector on the screen. The function **csco_ShowLoginForm('lform')** injects the logon form. **csco_ShowLanguageSelector('selector')** injects the Language Selector.

Modifying Your HTML File

Follow these steps to modify your HTML file:

Detailed Steps

Step 1 Name your file **logon.inc**. When you import the file, the ASA recognizes this filename as the logon screen.

Step 2 Modify the paths of images used by the file to include **/+CSCOU+/**.

Files that are displayed to remote users before authentication must reside in a specific area of the ASA cache memory represented by the path /+CSCOU+/. Therefore, the source for each image in the file must include this path. For example:

```
src="/+CSCOU+/asa5520.gif"
```

Step 3 Insert the special HTML code below. This code contains the Cisco functions, described earlier, that inject the login form and language selector onto the screen.

```
<body onload="csco_ShowLoginForm('lform');csco_ShowLanguageSelector('selector')">

<table>

<tr><td colspan=3 height=20 align=right><div id="selector" style="width:
300px"></div></td></tr>
<tr><td></td><td></td><td></td><td></td></tr>
<tr>
<td height="379"></td>
<td height="379"></td>
<td align=middle valign=middle>
<div id=lform >
<p> </p>
<p> </p>
<p> </p>
<p>Loading...</p>
</div>
</td>
</tr>
<tr>
<td width="251"></td>
<td width="1"></td>
```

```
<td align=right valign=right width="800">
<img border="1" src="/+CSCOU+/asa5500.jpg" width="660" height="220" align="middle">
</td></tr>

</table>
```

Customizing the Portal Page

Figure 91-26 shows the portal page and the pre-defined components you can customize:

Figure 91-26 Customizable Components of the Portal Page

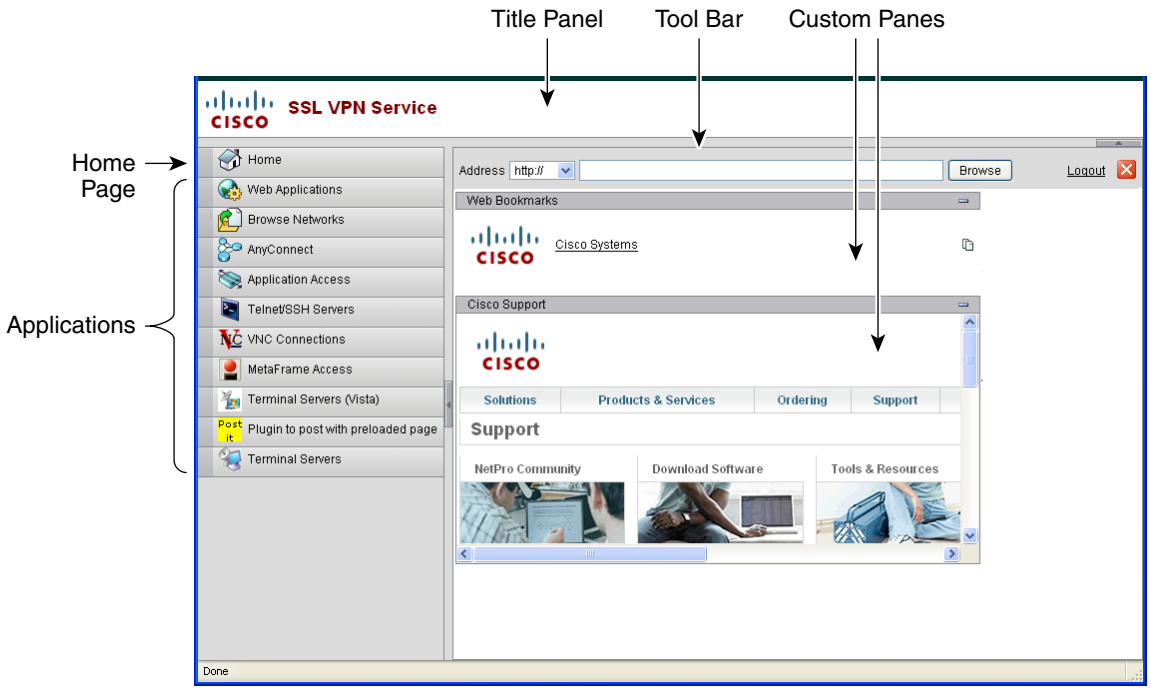

In addition to customizing the components of the page, you can divide the portal page into custom panes that display text, an image, an RSS feed, or HTML. In Figure 91-26, the portal page is divided into one column with two rows.

To customize the portal page, follow this procedure. You can preview your changes for each component by clicking the **Preview** button:

Step 1 Go to Portal Page and specify a title for the browser window.

Step 2 Display and customize the title panel. Go to Portal Page > Title Panel and check **Display title panel**. Enter text to display as the title and specify a logo. Specify any font styles.

Step 3 Enable and customize the toolbar. Go to Portal Page > Toolbar and check **Display toolbar**. Customize the Prompt Box, Browse button, and Logout prompt as desired.

Step 4 Customize the Applications list. Go to Portal Page > Applications and check **Show navigation panel**. The applications populated in the table are those applications you enabled in the ASA configuration, including client-server plugins and port forwarding applications.

Step 5 Create custom panes in the portal page space. Go to Portal Page > Custom Panes and divide the window into rows and columns for text, images, RSS feeds, or HTML pages, as desired.

Step 6 Specify a home page URL. Go to Portal Page > Home Page and check **Enable custom intranet web page**. Choose a bookmark mode that defines how bookmarks are organized.

Configure a timeout alert message and a tooltip. Go to Portal Page > Timeout Alerts. See Configuring Custom Portal Timeout Alerts for full instructions.

Configuring Custom Portal Timeout Alerts

So that users of the clientless SSL VPN feature can manage their time in the VPN session, the clientless SSL VPN portal page displays a countdown timer showing the total time left before the clientless VPN session expires. Sessions can timeout due to inactivity or because they have reached the end of a maximum allowed connection time that you have configured.

You can create custom messages to alert users that their session is about to end because of an idle timeout or a session timeout. Your custom message replaces the default idle timeout message. The default message is, "`Your session will expire in %s .`" The **`%s`** place holder in your message is replaced by a ticking countdown timer.

Step 1 Start ASDM and select **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Customization**.

Step 2 Click **Add** to add a new customization object or select an existing customization object and click **Edit** to add a custom idle timeout message to an existing customization object.

Step 3 In the Add / Edit Customization Object pane, expand the Portal Page node on the navigation tree and click **Timeout Alerts**.

Step 4 Check **Enable alert visual tooltip (red background for timer countdown)**. This displays the countdown timer as a tool tip on a red background. When users click the Time left area, the time area expands to display your custom timeout alert message. If you leave this box unchecked, users see the custom timeout alerts in a pop-up window.

Step 5 Enter a message in the Idle Timeout Message box and in the Session Timeout Message box. An example of a message could be, `Warning: Your session will end in %s. Please complete your work and prepare to close your applications.`

Step 6 Click **OK**.

Step 7 Click **Apply**.

Specifying a Custom Timeout Alert in a Customization Object File

If you desire, you can edit an existing customization object file outside of the ASA and import it to the ASA. For more information about Importing and Exporting Customization objects see Importing/Exporting Customization Object, page 91-94. See also, Creating XML-Based Portal Customization Objects and URL Lists, page 91-94.

The timeout messages are configured in the <timeout-alerts> XML element of your XML customization object file. The <timeout-alerts> element is a child of the <portal> element. The <portal> element is a child of the <custom> element.

The <timeout-alerts> element is placed after the <home-page> element and before any <application> elements in the order of the <portal> child elements.

You need to specify these child-elements of <timeout-alerts>:

- <alert-tooltip> – If set to “yes”, users see the countdown timer on a red background as a tool tip. Clicking the count down timer expands the tooltip to display your custom message. If set to “no” or if is undefined, users receive your custom messages in pop-up windows.
- <session-timeout-message> – Enter your custom session timeout message in this element. If set and not empty, users receive your custom message instead of the default message. The %s place holder in the message will be replaced with a ticking countdown timer.
- <idle-timeout-message> – Enter your custom idle timeout message in this element. If set and not empty, users receive your custom message instead of the default message. The %s place holder will be replaced with a ticking countdown timer.

Configuration Example for Timeout-alert Element and Child Elements

This example shows only the <timeout-alerts> elements of the <portal> element.

Note Do not cut and paste this example into an existing customization object.

```
<portal>
    <window></window>
    <title-panel></title-panel>
    <toolbar></toolbar>
    <url-lists></url-lists>
    <navigation-panel></navigation-panel>
    <home-page>
    <timeout-alerts>
        <alert-tooltip>yes</alert-tooltip>
        <idle-timeout-message>You session expires in %s due to idleness.</idle-timeout-message>
        <session-timeout-message>Your session expires in %s.</session-timeout-message>
    </timeout-alerts>
    <application></application>
    <column></column>
    <pane></pane>
    <external-portal></external-portal>
</portal>
```

Customizing the Logout Page

Figure 91-27 shows the logout page you can customize:

Figure 91-27 Components of the Logout Page

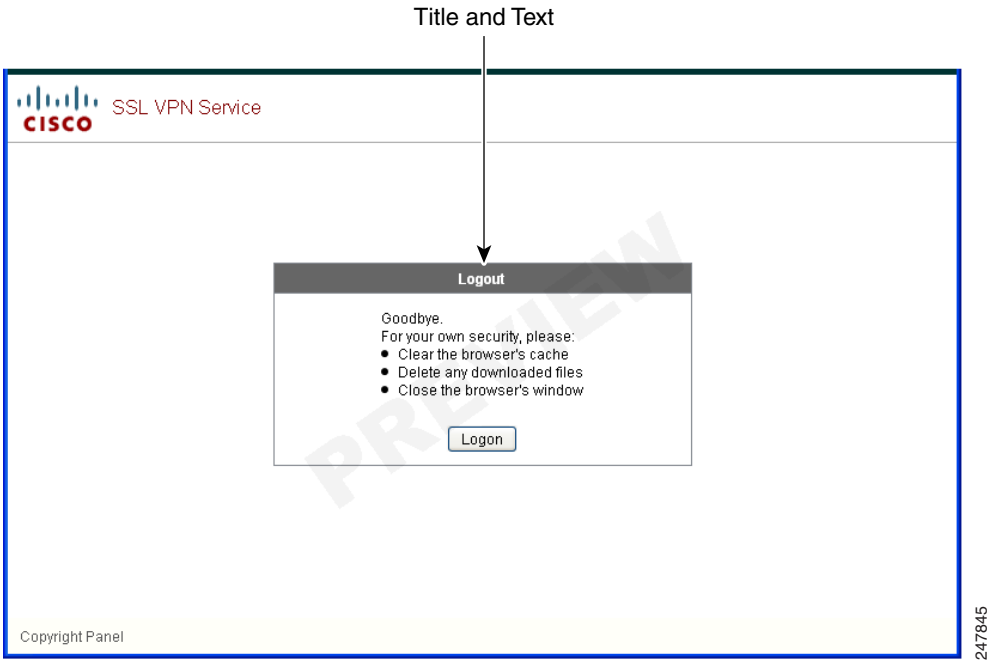

To customize the logout page, follow this procedure. You can preview your changes for each component by clicking the **Preview** button:

Step 1 Go to Logout Page. Customize the title or text as you desire.

Step 2 For the convenience of the user, you can display the Login button on the Logout page. To do this, check **Show logon button**. Customize the button text, if desired.

Step 3 Customize the title font or background, as desired.

Step 4 Click **OK**, then apply the changes to the customization object you edited.

Customizing the External Portal Page

Adding 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.

Detailed Steps

Step 1 Click **Add** and enter a name for the new customization object. Maximum 64 characters, no spaces.

Step 2 (Optional) Click **Find** to search for a customization object. Start typing in the field, and the tool searches the beginning characters of every field for a match. You can use wild cards to expand your search. For example, typing *sal* in the Find field matches a customization object named sales but not a customization object named wholesalers. If you type **sal* in the Find field, the search finds the first instance of either sales or wholesalers in the table.

Use the up and down arrows to skip up or down to the next string match. Check the **Match Case** checkbox to make your search case sensitive.

Step 3 Specify when the onscreen keyboard shows on portal pages. The choices are as follows:

- Do not show OnScreen Keyboard
- Show only for the login page
- Show for all portal pages requiring authentication

Step 4 (Optional) Highlight a customization object and click **Assign** to assign the selected object to one or more group policies, connection profiles, or LOCAL users.

Importing/Exporting 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 ASA for editing purposes, after which you can reimport it.

Detailed Steps

Step 1 Identify the customization object by name. Maximum 64 characters, no spaces.

Step 2 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 ASA.
- 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 ASA.
- Path—Identify the method to access the file (ftp, http, or https), and provide the path to the file.

Step 3 Click to import or export the file.

Creating XML-Based Portal Customization Objects and URL Lists

This section includes the following topics:

- Understanding the XML Customization File Structure, page 91-95

- Configuration Example for Customization, page 91-98
- Using the Customization Template, page 91-101
- Help Customization, page 91-113
- Import/Export Application Help Content, page 91-116

Understanding the XML Customization File Structure

Table 91-8 presents the file structure for an XML customization object.

Note 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 91-8 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	—
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	—

Table 91-8 XML-Based Customization File Structure (continued)

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
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	—

Table 91-8 XML-Based Customization File Structure (continued)

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	—
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	enableldisable	N/A	v
toolbar	**node**	—	—	—
mode	string	enableldisable	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.

Table 91-8 XML-Based Customization File Structure (continued)

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
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 url-mode	string string	— —	— —	URL for RSS,IMAGE or HTML type paned Modes: mangle, no-mangle
text	string	—	—	Text for TEXT type panes
column	number	—	—	—

Configuration Example for Customization

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>
  <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 ASA, 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 to create a customization object that meets the needs of your organization, import it to the ASA, 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 ASA.

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 ASA 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 91-9 shows the clientless application panels and predetermined filenames for the help content.

Table 91-9 Clientless Applications

Application Type	Panel	Filename
Standard	Application Access	app-access-hlp.inc
Standard	Browse Networks	file-access-hlp.inc

Table 91-9 Clientless Applications

Application Type	Panel	Filename
Standard	AnyConnect Client	net-access-hlp.inc
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[1]	ssh,telnet-hlp.inc
Plug-in	VNC Connections	vnc-hlp.inc

1. This plug-in is capable of doing both sshv1 and sshv2.

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 ASA 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

Detailed Steps

Step 1 Click **Import** to launch the Import Application Help Content dialog, where you can import new help content to flash memory for display during clientless sessions.

Step 2 (Optional) Click **Export** to retrieve previously imported help content selected from the table.

Step 3 (Optional) Click **Delete** to delete previously imported help content selected from the table.

Step 4 The abbreviation of the language rendered by the browser is displayed. 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**.

The filename that the help content file was imported as is provided.

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 ASA.

Step 2 Display the help file by appending the string in "URL of Help File in Flash Memory of the Security Appliance" in Table 91-10, to the address of the ASA, substituting *language* as described below, then press **Enter**.

Table 91-10 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**.

Note 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 91-11, 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.

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.

Detailed Steps

Step 1 The Language field specifies the language rendered by the browser but is not used for file translation. (This field is inactive in the Export Application Help Content dialog box.) Click the dots next to the Language field and double-click the row containing the language shown in the Browse Language Code dialog box. Confirm the abbreviation in the Language Code field matches the abbreviation in the row and click **OK**.

Step 2 If the language for which you want to provide help content is not present in the Browse Language Code dialog box, perform the following

1. Display the list of languages and abbreviations rendered by your browser.
2. Enter the abbreviation for the language in the Language Code field and click **OK.**

   OR

   You can also enter it into the Language text box to the left of the dots.

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**.

Step 3 If you are importing, choose the new help content file from the File Name drop-down list. If you are exporting, this field is unavailable.

Step 4 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 ASA:
  - 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.

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 ASA.

Step 2 Display the help file by appending the string in "URL of Help File in Flash Memory of the Security Appliance" in Table 91-11, to the address of the ASA, substituting *language* as described below, then press **Enter**.

Table 91-11 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**.

Note 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 91-11, 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 91-13.

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.

Configuring Browser Access to Client-Server Plug-ins

The Client-Server Plug-in table displays the plug-ins the ASA 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
- Preparing the Security Appliance for a Plug-in
- Installing Plug-ins Redistributed By Cisco

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 ASA 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.

The ASA 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 ASA 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 91-12 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 91-12 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 ASA obtains the plug-ins from the primary ASA.

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 ASA.

Before installing the first plug-in, you must follow the instructions in the next section.

Prerequisites

- The plug-ins do not work if the security appliance configures the clientless session to use a proxy server.

Note The remote desktop protocol plug-in does not support load balancing with a session broker. Because of the way the protocol handles the redirect from the session broker, the connection fails. If a session broker is not used, the plug-in works.

- The plug-ins support single sign-on (SSO). They use the *same* credentials entered to open the clientless SSL VPN session. Because the plug-ins do not support macro substitution, you do not have the options to perform SSO on different fields such as the internal domain password or on an attribute on a RADIUS or LDAP server.
- To configure SSO support for a plug-in, you install the plug-in, add a bookmark entry to display a link to the server, and specify SSO support when adding the bookmark.
- The minimum access rights required for remote use belong to the guest privilege mode.

Requirements

- 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.

- Clientless SSL VPN must be enabled on the ASA to provide remote access to the plug-ins.
- A stateful failover does not retain sessions established using plug-ins. Users must reconnect following a failover.
- Plug-ins require that ActiveX or Oracle Java Runtime Environment (JRE) 1.4.2 (or later) is enabled on the browser. There is no ActiveX version of the RDP plug-in for 64-bit browsers.

RDP Plug-in ActiveX Debug Quick Reference

To set up and use an RDP plug-in, you must add a new environment variable. For the process of adding a new environment variable, use the following steps:

Step 1 Right-click **My Computer** to access the System Properties, and choose the **Advanced** tab.

Step 2 On the Advanced tab, choose the environment variables button.

Step 3 In the new user variable dialog box, enter the RF_DEBUG variable.

Step 4 Verify the new Environment Variable in the user variables section.

Step 5 If you used the client computer with versions of WebVPN before version 8.3, you must remove the old Cisco Portforwarder Control. Go to the C:/WINDOWS/Downloaded Program Files directory, right-click portforwarder control, and choose **Remove**.

Step 6 Clear all of the Internet Explorer browser cache.

Step 7 Launch your WebVPN session and establish an RDP session with the RDP ActiveX Plug-in.

You can now observe events in the Windows Application Event viewer.

Preparing the Security Appliance for a Plug-in

Before installing a plug-in, prepare the ASA by performing the following steps:

Step 1 Make sure clientless SSL VPN (“webvpn”) is enabled on an ASA interface.

Step 2 Install an SSL certificate onto the ASA 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 ASA. The remote PC must be able to use DNS or an entry in the System32\drivers\etc\hosts file to resolve the FQDN.

Customizing Help

The ASA displays help content on the application panels during VPN sessions. You can customize the help files provided by Cisco or create help files in other languages. You then import them to flash memory for display during subsequent sessions. You can also retrieve previously imported help content files, modify them, and reimport them to flash memory.

Each application panel displays its own help file content using a predetermined filename. The prospective location of each is in the /+CSCOE+/help/*language*/ URL within flash memory of the ASA. Table 91-13 shows the details about each of the help files you can maintain for VPN sessions.

Table 91-13 VPN Application Help Files

Application Type	Panel	URL of Help File in Flash Memory of the Security Appliance	Help File Provided By Cisco in English?
Standard	Application Access	/+CSCOE+/help/*language*/app-access-hlp.inc	Yes
Standard	Browse Networks	/+CSCOE+/help/*language*/file-access-hlp.inc	Yes
Standard	AnyConnect Client	/+CSCOE+/help/*language*/net-access-hlp.inc	Yes
Standard	Web Access	/+CSCOE+/help/*language*/web-access-hlp.inc	Yes
Plug-in	MetaFrame Access	/+CSCOE+/help/*language*/ica-hlp.inc	No
Plug-in	Terminal Servers	/+CSCOE+/help/*language*/rdp-hlp.inc	Yes
Plug-in	Telnet/SSH Servers	/+CSCOE+/help/*language*/ssh,telnet-hlp.inc	Yes
Plug-in	VNC Connections	/+CSCOE+/help/*language*/vnc-hlp.inc	Yes

language is 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 specify a particular language code, copy the language abbreviation from 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**.

The following sections describe how to customize the help contents:

- Customizing a Help File Provided By Cisco, page 91-121
- Creating Help Files for Languages Not Provided by Cisco, page 91-122
- Requiring Usernames and Passwords, page 91-122

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 as follows:

Detailed Steps

Step 1 Use your browser to establish a clientless SSL VPN session with the ASA.

Step 2 Display the help file by appending the string in "URL of Help File in Flash Memory of the Security Appliance" in Table 91-13, to the address of the ASA, then press Enter.

Note Enter **en** in place of *language* to get the help file in English.

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**.

Note 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 modify the file.

Note You can use most HTML tags, but do *not* use tags that define the document and its structure (e.g., 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 91-13, and that it does not have an extra filename extension.

See "Requiring Usernames and Passwords" to import the modified file.

Creating Help Files for Languages Not Provided by Cisco

Use HTML to create help files in other languages.

We recommend creating a separate folder for each language you want to support.

Save the file as HTML only. Use the filename following the last slash in "URL of Help File in Flash Memory of the Security Appliance" in Table 91-13.

See the next section to import the files for display during VPN sessions.

Restrictions

You can use most HTML tags, but do *not* use tags that define the document and its structure (e.g., 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.

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 91-14 lists the type of usernames and passwords that clientless SSL VPN users might need to know.

Table 91-14 Usernames and Passwords to Give to Users of Clientless SSL VPN Sessions

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 clientless SSL VPN
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 to always click the logout icon on the toolbar to close the clientless SSL VPN session. (Closing the browser window does not close the session.)

Clientless SSL VPN ensures the security of data transmission between the remote PC or workstation and the ASA on the corporate network. Advise users that using clientless SSL VPN does not ensure that communication with every site is secure. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate ASA to the destination web server is not private because it is not encrypted.

"Observing Clientless SSL VPN Security Precautions" on page 5 addresses an additional tip to communicate with users, depending on the steps you follow within that section.

Configuring Remote Systems to Use Clientless SSL VPN Features

This section describes how to set up remote systems to use clientless SSL VPN and includes the following topics:

- Starting Clientless SSL VPN, page 91-124
- Using the Clientless SSL VPN Floating Toolbar, page 91-124
- Browsing the Web, page 91-125
- Browsing the Network (File Management), page 91-125
- Using Port Forwarding, page 91-127
- Using E-mail Via Port Forwarding, page 91-128

- Using E-mail Via Web Access, page 91-129
- Using E-mail Via E-mail Proxy, page 91-129
- Using Smart Tunnel, page 91-129

You may configure user accounts differently and different clientless SSL VPN features can be available to each user.

Starting Clientless SSL VPN

You can connect to the internet using any supported connection including:

- home DSL, cable, or dial-ups
- public kiosks
- hotel hook-ups
- airport wireless nodes
- internet cafes

Note See the *Cisco ASA 5500 Series VPN Compatibility Reference* for the list of web browsers supported by clientless SSL VPN.

Prerequisites

- Cookies must be enabled on the browser in order to access applications via port forwarding.
- You must have a URL for clientless SSL VPN. The URL must be an https address in the following form: https://*address*, where *address* is the IP address or DNS hostname of an interface of the ASA (or load balancing cluster) on which SSL VPN is enabled. For example, https://cisco.example.com.
- You must have a clientless SSL VPN username and password.

Restrictions

- Clientless SSL VPN supports local printing, but it does not support printing through the VPN to a printer on the corporate network.

Using the Clientless SSL VPN Floating Toolbar

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.

The floating toolbar represents the current clientless SSL VPN session. If you click the **Close** button, the ASA prompts you to confirm that you want to close the clientless SSL VPN session.

Tip To paste text into a text field, use Ctrl-V. (Right-clicking is disabled on the toolbar displayed during the clientless SSL VPN session.)

Restrictions

If you configure your browser to block popups, the floating toolbar cannot display.

Browsing the Web

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 title bar for clientless SSL VPN 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 websites that appear as links on the clientless SSL VPN Home page are available

Prerequisites

- You need the username and password for protected websites.

Restrictions

Also, depending on how you configured a particular account, it might be that:

- Some websites are blocked
- Only the websites that appear as links on the clientless SSL VPN Home page are available

Browsing the Network (File Management)

Users might not be familiar with how to locate their files through your organization network.

Note 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.

Prerequisites

- You must configure file permissions for shared remote access.
- You must have the server names and passwords for protected file servers.
- You must have the domain, workgroup, and server names where folders and files reside.

Restrictions

Only shared folders and files are accessible via clientless SSL VPN.

Using the Remote File Explorer

The Remote File Explorer provides the user with a way to browse the corporate network from their web browser. When the users clicks the Remote File System icon on the Cisco SSL VPN portal page, an applet is launched on the user's system displaying the remote file system in a tree and folder view.

Figure 91-28 Clientless SSL VPN Remote File Explorer

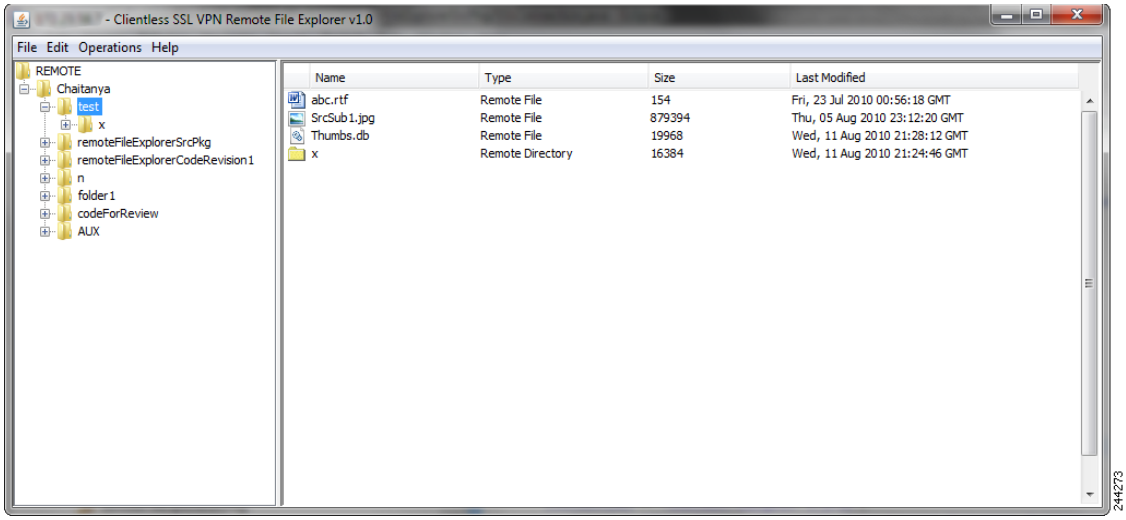

The browser enables the user to:

- Browse the remote file system
- Rename files
- Move or copy files within the remote file system and between the remote and local file systems.
- Perform bulk uploads and downloads of files

Note This functionality requires the Oracle Java Runtime Environment (JRE) 1.4 or later is installed on the user's machine and Java enabled in the web browser. Launching remote files requires JRE 1.6 or later.

Renaming a File or Folder

To rename a file or folder:

Step 1 Click the file or folder to be renamed.

Step 2 Select **Edit** > **Rename**.

Step 3 When prompted, enter the new name in the dialog.

Step 4 Click **OK** to rename the file or folder. Alternative, click **Cancel** to leave the name unchanged.

Moving or Copying Files or Folders on the Remote Server

To move or copy a file or folder on the remote server:

Step 1 Navigate to the source folder containing the file or folder to be moved or copied.

Step 2 Click the file or folder.

Step 3 To copy the file select **Edit** > **Copy**. Alternatively, to move the file select **Edit** > **Cut**.

Step 4 Navigate to the destination folder.

Step 5 Select **Edit** > **Paste**.

Copying Files from the Local System Drive to the Remote Folder

You can copy files between the local file system and the remote file system by dragging and dropping them between the right pane of the Remote File Browser and your local file manager application.

Uploading and Downloading Files

You can download a file by clicking it in the browser, selecting **Operations** > **Download**, and providing a location and name to save the file in the **Save** dialog.

You can upload a file by clicking the destination folder, selecting **Operations** > **Upload**, and providing the location and name of the file in the **Open** dialog,

This functionality has the following restrictions:

- The user cannot view sub-folders for which they are not permitted access.
- Files that the user is not permitted to access cannot be moved or copied, even though they are displayed in the browser.
- The maximum depth of nested folders is 32.
- The tree view does not support drag and drop copying.
- When moving files between multiple instances of the Remote File Explorer, all instances must be exploring the same server (root share).
- The Remote File Explorer can display a maximum of 1500 files and folders in a single folder. If a folder exceeds this limit the folder cannot be displayed.

Using Port Forwarding

Note Users should always close the Application Access window when they finish using applications by clicking the **Close** icon. Failure to quit the window properly can cause Application Access or the applications themselves to be disabled. See Recovering from hosts File Errors When Using Application Access for details.

Prerequisites

- On Mac OS X, only the Safari browser supports this feature.
- You must have client applications installed.
- You must have Cookies enabled on the browser.
- You must have administrator access on the PC if you use DNS names to specify servers, because modifying the hosts file requires it.
- You must have Oracle Java Runtime Environment (JRE) version 1.4.x and 1.5.x installed.

  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:

  a. Clear the browser cache and close the browser.

b. Verify that no Java icons are in the computer task bar.

c. Close all instances of Java.

d. Establish a clientless SSL VPN session and launch the port forwarding Java applet.

- You must have JavaScript enabled on the browser. By default, it is enabled.
- If necessary, you must configure client applications.

Note The Microsoft Outlook client does not require this configuration step. All non-Windows client applications require configuration. To determine if configuration is necessary for a Windows application, check the value of the Remote Server field. If the Remote Server field 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.

Restrictions

Because this feature requires installing Oracle Java Runtime Environment (JRE) and configuring the local clients, and because doing so requires administrator permissions on the local system or full control of C:\windows\System32\drivers\etc, it is unlikely that users will be able to use applications when they connect from public remote systems.

Detailed Steps

To configure the client application, use the server's locally mapped IP address and port number. To find this information:

1. Start a clientless SSL VPN session and click the **Application Access** link on the 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 a clientless SSL VPN session does not open the site over that session. To open a site over the session, paste the URL into the Enter Clientless SSL VPN (URL) Address field.

Using E-mail Via Port Forwarding

To use e-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.

Prerequisites

You must fulfill requirements for application access and other mail clients.

Restrictions

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

The following e-mail applications are supported:

- Microsoft Outlook Web App to Exchange Server 2010.

  OWA requires Internet Explorer 7 or later, or Firefox 3.01 or later.
- Microsoft Outlook Web Access to Exchange Server 2007, 2003, and 2000.

  For best results, use OWA on Internet Explorer 8.x or later, or Firefox 8.x.
- Lotus iNotes

Prerequisites

You must have the web-based e-mail product installed.

Restrictions

Other web-based e-mail applications should also work, but we have not verified them.

Using E-mail Via E-mail Proxy

The following legacy e-mail applications are supported:

- Microsoft Outlook 2000 and 2002
- Microsoft Outlook Express 5.5 and 6.0

See the instructions and examples for your mail application in Using E-Mail over Clientless SSL VPN.

Prerequisites

- You must have the SSL-enabled mail application installed.
- Do not set the ASA SSL version to TLSv1 Only. Outlook and Outlook Express do not support TLS.
- You must have your mail application properly configured.

Restrictions

- Other SSL-enabled clients should also work, but we have not verified them.

Using Smart Tunnel

Administration privileges are not required to use Smart Tunnel.

Note Java is not automatically downloaded for you as in port forwarder.

Prerequisites

- Smart tunnel requires either ActiveX or JRE (1.4x and 1.5x) on Windows and Java Web Start on Mac OS X.
- You must ensure cookies enabled on the browser.
- You must ensure JavaScript is enabled on the browser.

Restrictions

- Mac OS X does not support a front-side proxy.
- Supports only the operating systems and browsers specified in "Configuring Smart Tunnel Access" section on page 91-45.
- Only TCP socket-based applications are supported.

The ASA includes a translation table template for each domain that is part of 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 in this file are empty. You can edit the messages and import the template to create a new translation table object that resides in flash memory.

You can also export an existing translation 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 translation table object, overwriting previous messages.

Some templates are static, but some change based on the configuration of the ASA. Because you can customize the *logon and logout pages, portal page, and URL bookmarks for clientless users,* the **ASA generates the customization** and **url-list** translation domain templates dynamically, and the template automatically reflects your changes to these functional areas.

After creating translation tables, they are available to customization objects that you create and apply to group policies or user attributes. With the exception of the AnyConnect translation domain, a translation table has no affect, and messages are not translated on user screens until you create a customization object, identify a translation table to use in that object, and specify that customization for the group policy or user. Changes to the translation table for the AnyConnect domain are immediately visible to AnyConnect client users.

Note

Step 1 Navigate to **Configuration > Remote Access VPN > 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.

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.

Adding/Editing 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.

Step 1 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.

Step 2 Specify a language. Use an abbreviation that is compatible with the language options of your browser. The ASA creates the new translation table with this name.

Step 3 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.

Customizing the AnyConnect Client

You can customize the AnyConnect VPN client to display your own corporate image to remote users, including clients running on Windows, Linux, and Mac OS X computers.

You can use one of three methods to customize the client:

- Rebrand the client by importing individual client GUI components, such as the corporate logo and icons, to the ASA which deploys them to remote computers with the installer.
- Import your own program (Windows and Linux only) that provides its own GUI or CLI and uses the AnyConnect API.
- Import a transform (Windows only) that you create for more extensive rebranding. The ASA deploys it with installer.
- Create Scripts that deploy with the client and run when the client establishes or terminates a VPN connection.

The following sections explain how to customize the AnyConnect client:

- Customizing AnyConnect by Importing Resource Files, page 91-132
- Customizing Your Own AnyConnect GUI Text and Scripts, page 91-133
- Customizing AnyConnect GUI Text and Messages, page 91-136
- Customizing the Installer Program Using Installer Transforms, page 91-137
- Localizing the Install Program using Installer Transforms, page 91-138

Restrictions

- Customization is not supported for the AnyConnect client running on a Windows Mobile device.

Customizing AnyConnect by Importing Resource Files

You can customize the AnyConnect client by importing your own custom files to the security appliance, which deploys the new files with the client. For detailed information about the original GUI icons and information about their sizes, see the *AnyConnect VPN Client Administrators Guide*. You can use this information to create your custom files.

Detailed Steps

To import and deploy your custom files with the client, follow this procedure:

Step 1 Go to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Customization/Localization > Resources**.

Click **Import**. The Import AnyConnect Customization Object window displays.

Step 2 Enter the Name of the file to import. See the *AnyConnect VPN Client Administrators Guide* for the filenames of all the GUI components that you can replace.

Note The filenames of your custom components must match the filenames used by the AnyConnect client GUI. The filenames of the GUI components are different for each OS and are case sensitive for Mac and Linux. For example, if you want to replace the corporate logo for Windows clients, you must import your corporate logo as *company_logo.bmp*. If you import it as a different filename, the AnyConnect installer does not change the component. However, if you deploy your own executable to customize the GUI, the executable can call resource files using any filename.

Step 3 Select a platform and specify the file to import. Click **Import Now**. The file now appears in the table.

Note If you import an image as a resource file (such as company_logo.bmp), the image you import customizes the AnyConnect client until you reimport another image using the same filename. For example, if you replace company_logo.bmp with a custom image, and then delete the image, the client continues to display your image until you import a new image (or the original Cisco logo image) using the same filename.

Step 4 Clicking **Import** launches the Import AnyConnect Customization Objects dialog, where you can specify a file to import as an object.

Step 5 Clicking **Export** launches the Export AnyConnect Customization Objects dialog, where you can specify a file to export as an object.

Step 6 Clicking **Delete** removes the selected object.

The type of remote PC platform supported by the object and the object name is displayed.

Customizing Your Own AnyConnect GUI Text and Scripts

For Windows, Linux, or Mac (PPP or Intel-based) computers, you can deploy your own client that uses the AnyConnect client API. You replace the AnyConnect GUI or the AnyConnect CLI by replacing the client binary files.

You can also download and run scripts that run when the client establishes a connection (an *OnConnect* script), or when the client terminates a session (an *OnDisconnect* script). Some examples that show how you might want to use this feature include:

- Refreshing the group policy upon VPN connection.
- Mapping a network drive upon VPN connection, and un-mapping it after disconnection.
- Logging on to a service upon VPN connection, and logging off after disconnection.

For complete information about customizing the AnyConnect GUI and creating and deploying scripts, see the *AnyConnect VPN Client Administrators Guide*.

The following sections describe how to import binary executables and scripts to the ASA:

- Importing your own GUI as a Binary Executable, page 91-133
- Importing Scripts, page 91-134

Importing your own GUI as a Binary Executable

For Windows, Linux, or Mac (PPP or Intel-based) computers, you can deploy your own client that uses the AnyConnect client API. You replace the AnyConnect GUI or the AnyConnect CLI by replacing the client binary files. Table 91-15 lists the filenames of the client executable files for the different operating systems.

Table 91-15 Filenames of Client Executables

Client OS	Client GUI File	Client CLI File
Windows	vpnui.exe	vpncli.exe

Table 91-15 Filenames of Client Executables

Client OS	Client GUI File	Client CLI File
Linux	vpnui	vpn
Mac	Not supported[1]	vpn

1. Not supported by ASA deployment. However, you can deploy an executable for the Mac that replaces the client GUI using other means, such as Altiris Agent.

Your executable can call any resource files, such as logo images, that you import to the ASA (See Table 91-15). Unlike replacing the pre-defined GUI components, when you deploy your own executable, can use any filenames for your resource files.

We recommend that you sign your custom Windows client binaries (either GUI or CLI version) that you import to the ASA. A signed binary has a wider range of functionality available to it. If the binaries are not signed the following functionality is affected:

- Web-Launch—The clientless portal is available and the user can authenticate. However, the behavior surrounding tunnel establishment does not work as expected. Having an unsigned GUI on the client results in the client not starting as part of the clientless connection attempt. And once it detects this condition, it aborts the connection attempt.
- SBL—The Start Before Logon feature requires that the client GUI used to prompt for user credentials be signed. If it is not, the GUI does not start. Because SBL is not supported for the CLI program, this affects only the GUI binary file.
- Auto Upgrade—During the upgrade to a newer version of the client, the old GUI exits, and after the new GUI installs, the new GUI starts. The new GUI does not start unless it is signed. As with Web-launch, the VPN connection terminates if the GUI is not signed. However, the upgraded client remains installed.

Restrictions

The ASA 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*.

Importing Scripts

AnyConnect lets you download and run scripts when the following events occur:

- Upon the establishment of a new AnyConnect client VPN session with the security appliance. We refer to a script triggered by this event as an *OnConnect* script because it requires this filename prefix.
- Upon the tear-down of an AnyConnect client VPN session with the security appliance. We refer to a script triggered by this event as an *OnDisconnect* script because it requires this filename prefix.

Thus, the establishment of a new AnyConnect VPN session initiated by Trusted Network Detection triggers the OnConnect script (assuming the requirements are satisfied to run the script). The reconnection of a persistent AnyConnect VPN session after a network disruption does not trigger the OnConnect script.

Prerequisites

These instructions assume you know how to write scripts and run them from the command line of the targeted endpoint to test them.

Restrictions

- The AnyConnect software download site provides some example scripts; if you examine them, please remember that they are only examples; they may not satisfy the local computer requirements for running them, and are unlikely to be usable without customizing them for your network and user needs. Cisco does not support example scripts or customer-written scripts.

For complete information about deploying scripts, and their limitations and restrictions, see the *AnyConnect VPN Client Administrators Guide*.

Writing, Testing, and Deploying Scripts

Deploy AnyConnect scripts as follows:

Restrictions

- Scripts written on Microsoft Windows computers have different line endings than scripts written on Mac OS X and Linux. Therefore, you should write and test the script on the targeted OS. If a script cannot run properly from the command line on the native OS, AnyConnect cannot run it properly either.
- Microsoft Windows Mobile does not support this option. You must deploy scripts using the manual method for this OS.

Step 1 Write and test the script using the OS type on which it will run when AnyConnect launches it.

Step 2 To import a script, go to **Network (Client) Access > AnyConnect Customization/Localization > Script.** The Customization Scripts pane displays.

Step 3 Enter a name for the script. Be sure to specify the correct extension with the name. For example, *myscript*.bat.

Step 4 Choose a script action: *Script runs when client connects* or *Script runs when client disconnects*.

AnyConnect adds the prefix *scripts_* and the prefix *OnConnect* or *OnDisconnect* to your filename to identify the file as a script on the ASA. When the client connects, the ASA downloads the script to the proper target directory on the remote computer, removing the *scripts_* prefix and leaving the remaining *OnConnect* or *OnDisconnect* prefix. For example, if you import the script *myscript.bat*, the script appears on the ASA as *scripts_OnConnect_myscript.bat*. On the remote computer, the script appears as *OnConnect_myscript.bat.*

To ensure the scripts run reliably, configure all ASAs to deploy the same scripts. If you want to modify or replace a script, use the same name as the previous version and assign the replacement script to all of the ASAs that the users might connect to. When the user connects, the new script overwrites the one with the same name.

Step 5 Select a file as the source of the script. The name does not need to be the same as the name you provided for the script. ASDM imports the file from any source file, creating the new name you specify for Name in Step 3.

Table 91-16 shows the locations of scripts on the remote computer:

Table 91-16 Required Script Locations

OS	Directory
Microsoft Windows 7 and Vista	%ALLUSERPROFILE%\Cisco\Cisco AnyConnect VPN Client\Scripts
Microsoft Windows XP	%ALLUSERPROFILE%\Application Data\Cisco\Cisco AnyConnect VPN Client\Scripts

Table 91-16 Required Script Locations

OS	Directory
Linux	/opt/cisco/vpn/scripts **Note** Assign execute permissions to the file for User, Group and Other.
Mac OS X	/opt/cisco/vpn/scripts
Windows Mobile	%PROGRAMFILES%\Cisco AnyConnect VPN Client\Scripts

Step 6 Click **Import** to launch the Import AnyConnect Customization Objects dialog, where you can specify a file to import as an object.

Step 7 Click **Export** to launch the Export AnyConnect Customization Objects dialog, where you can specify a file to export as an object.

Step 8 Click **Delete** to remove a selected object.

The type of remote PC platform supported by the object and the object name is displayed.

Customizing AnyConnect 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. The text contained between the quotes of the msgid field represents the default text. *Do not* change this text. To display a different message, insert your custom text between the quotes of mgstr. 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."
msgstr ""
```

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** ASA.

Step 6 (Optional) Click **Add** to launch the Add Localization Entry dialog where you can select a localization template to add and you can edit the contents of the template.

Step 7 (Optional) Click **Edit** to launch the Edit Localization Entry dialog for the selected language in the table, and allows you to edit the previously-imported language localization table.

Step 8 (Optional) Click **Delete** to delete a selected language localization table.

Step 9 (Optional) Click **Import** to launch the Import Language Localization dialog where you can import a language localization template or table.

Step 10 (Optional) Click **Export** to launch 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.

Step 11 (Optional) Specify the language of the localization table.

Customizing the Installer Program Using Installer Transforms

You can perform more extensive customizing of the AnyConnect client GUI (Windows only) by creating your own transform that deploys with the client installer program. You import the transform to the ASA, which deploys it with the installer program.

To create an MSI transform, you can download and install the free database editor from Microsoft, named Orca. With this tool, you can modify existing installations and even add new files. The Orca tool is part of the Microsoft Windows Installer Software Development Kit (SDK) which is included in the Microsoft Windows SDK. The following link leads to the bundle containing the Orca program:

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/msi/setup/orca_exe.asp.

After you install the SDK, the Orca MSI is located here:

C:\Program Files\Microsoft SDK SP1\Microsoft Platform SDK\Bin\Orca.msi.

Install the Orca software, then access the Orca program from your Start > All Programs menu.

To import your transform, follow these steps:

Step 1 Go to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Customization/Localization > Customized Installer Transforms**. Click **Import**. The Import AnyConnect Customization Objects windows displays.

Step 2 Enter the Name of the file to import. Unlike the names of other customizing objects, the name is not significant to the ASA and is for your own convenience.

Step 3 Select a platform and specify the file to import. Click **Import Now**. The file now appears in the table.

Note Windows is the only valid choice for applying a transform.

Step 4 (Optional) Click **Import** to launch the Import AnyConnect Customization Objects dialog, where you can specify a transform file to import.

Step 5 (Optional) Click **Export** to launch the Export AnyConnect Customization Objects dialog, where you can specify a transform file to export.

Step 6 (Optional) Click **Delete** to remove the selected file.

The type of remote PC platform supported by the transform and the name of the transform is displayed.

Configuration Example for Transform

While offering a tutorial on creating transforms is beyond the scope of this document, we provide the text below as representative of some entries in a transform. These entries replace *company_logo.bmp* with a local copy and install the custom profile *MyProfile.xml*.

```
DATA CHANGE - Component Component ComponentId
+ MyProfile.xml {39057042-16A2-4034-87C0-8330104D8180}

Directory_ Attributes Condition KeyPath
Profile_DIR 0 MyProfile.xml

DATA CHANGE - FeatureComponents Feature_ Component_
  + MainFeature MyProfile.xml

DATA CHANGE - File File Component_ FileName FileSize Version Language Attributes Sequence
+ MyProfile.xml MyProfile.xml MyProf~1.xml|MyProfile.xml 601 8192 35
 <> company_logo.bmp 37302{39430} 8192{0}

DATA CHANGE - Media DiskId LastSequence DiskPrompt Cabinet VolumeLabel Source
  + 2 35
```

Specify transform files for customizing the AnyConnect client installation in this pane.

Localizing the Install Program using Installer Transforms

As with the AnyConnect client GUI, you can translate messages displayed by the client installer program. The ASA uses transforms to translate the messages displayed by the installer. The transform alters the installation, but leaves the original security-signed MSI intact. These transforms only translate the installer screens and do not translate the client GUI screens.

Each language has its own transform. You can edit a transform with a transform editor such as Orca, and make changes to the message strings. Then you import the transform to the ASA. When the user downloads the client, the client detects the preferred language of the computer (the locale specified during installation of the operating system) and applies the appropriate transform.

We currently offer transforms for 30 languages. These transforms are available in the following .zip file on the AnyConnect client software download page at cisco.com:

anyconnect-win-<VERSION>-web-deploy-k9-lang.zip

In this file, `<VERSION>` is the version of AnyConnect release (e.g. 2.2.103).

The package contains the transforms (.mst files) for the available translations. If you need to provide a language to remote users that is not one of the 30 languages we provide, you can create your own transform and import it to the ASA as a new language. With Orca, the database editor from Microsoft, you can modify existing installations and new files. Orca is part of the Microsoft Windows Installer Software Development Kit (SDK) which is included in the Microsoft Windows SDK. The following link leads to the bundle containing the Orca program:

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/msi/setup/orca_exe.asp.

After you install the SDK, the Orca MSI is located here:

C:\Program Files\Microsoft SDK SP1\Microsoft Platform SDK\Bin\Orca.msi.

The following procedure shows how to import a transform to the ASA using ASDM:

Step 1 Import a Transform. Go to: **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Customization/Localization > Localized Installer Transforms**. Click **Import**. The Import MST Language Localization window opens.

Step 2 Choose a language for this transform. Click the Language drop-list to display languages and their industry-recognized abbreviations. If you enter the abbreviation manually, be sure to use an abbreviation recognized by browsers and operating systems.

Step 3 Click **Import Now**. A message displays saying you successfully imported the table.

Be sure to click **Apply** to save your changes.

Step 4 (Optional) Click **Import** to launch the Import AnyConnect Customization Objects dialog, where you can specify a file to import as a transform.

Step 5 (Optional) Click **Export** to launch the Export AnyConnect Customization Objects dialog, where you can specify a file to export as a transform.

Step 6 (Optional) Click **Delete** to remove the selected transform.

The type of remote PC platform supported by the transform and the name are displayed.

Importing/Exporting Language Localization

In the Import Translation Table and Export Translation Table dialog boxes, you can import or export a translation table to the ASA 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.

Detailed Steps

Step 1 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.

Step 2 The name of the XML file containing the message fields includes the following:

- 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. This plug-in is capable of doing both sshv1 and sshv2.
- plugin-vnc—Messages for the VNC plug-in.

Step 3 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 ASA.
- Path—Identify the method to access the file (ftp, http, or https), and provide the file location.
- Flash file system—Choose this method to export a file that resides on the ASA.
- Path—File location.
- 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.

Step 4 Click **Import/Export Now** to import or export the file.

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.

You can now use bookmarks with macro substitutions for auto sign-on on some web pages. The former POST plug-in approach was created so that administrators could specify a POST bookmark with sign-on macros and receive a kick-off page to load prior to posting the POST request. This POST plug-in approach eliminated those requests that required the presence of cookies or other header items. Now an an administrator determines the pre-load page and URL, which specifies where you want the post login request sent. A pre-load page enables an endpoint browser to fetch certain information that is sent along to the webserver or web application rather than just using a POST request with credentials.

The existing bookmark lists are displayed. You can add, edit, delete, import, or export the bookmark list. You can configure lists of servers and URLs for access and order the items in the designated URL list.

Guidelines

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 the "Observing Clientless SSL VPN Security Precautions" section on page 91-5 for instructions.

Detailed Steps

Step 1 Specify the name of the list to be added or select the name of the list to be modified or deleted.

The bookmark title and actual associated URL are displayed.

Step 2 (Optional) Click **Add** to configure a new server or URL. See these procedures for additional information:

- Adding a Bookmark for a URL with a GET or Post Method, page 91-141
- Adding a URL for a Predefined Application Template, page 91-142
- Adding a Bookmark for an Auto Sign-On Application, page 91-144

Step 3 (Optional) Click **Edit** to make changes to the server, URL, or display name.

Step 4 (Optional) Click **Delete** to remove the selected item from the URL list. No confirmation or undo exists.

Step 5 (Optional) 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 ASA.
- Remote server—Click to import a file that resides on a remote server accessible from the ASA.
- 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.

Step 6 (Optional) Highlight a bookmark and click **Assign** to assign the selected bookmark to one or more group policies, dynamic access policies, or LOCAL users.

Step 7 (Optional) Change the position of the selected item in the URL list using the **Move Up** or **Move Down** options.

Step 8 Click OK.

Adding a Bookmark for a URL with a GET or Post Method

The Add Bookmark Entry dialog box lets you create a link or bookmark for a URL list.

Prerequisites

To access a shared folder on your network, use the format \\server\share\subfolder\<*personal folder*>. The user must have list permission for all points above <*personal folder*>.

Detailed Steps

Step 1 Navigate to **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks**, and click the **Add** button.

Step 2 Select **URL with GET or POST method** to use for bookmark creation.

Step 3 Enter a name for thise bookmark, which will be displayed on the portal.

Step 4 Use the URL drop-down menu to select the URL type: http, https, cifs, or ftp. The URL drop-down shows standard URL types, plus types for all the plug-ins you installed.

Step 5 Enter the DNS name or IP address for this bookmark (URL). 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**

Step 6 (Optional) Enter a preload URL. When you enter a preload URL, you can also enter the wait time, which is the time you allow for loading of the page until you are forwarded to the actual POST URL.

Step 7 As a subtitle, provide additional user-visible text that describes the bookmark entry.

Step 8 Use the Thumbnail drop-down menu to select an icon to associate with the bookmark on the end-user portal.

Step 9 Click **Manage** to import or export images to use as thumbnails.

Step 10 Click to open the bookmark in a new window that uses the smart tunnel feature to pass data through the ASA 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.

Step 11 Check **Allow the users to bookmark the link** to let clientless SSL VPN users use the Bookmarks or Favorites options on their browsers. Uncheck to prevent access to these options. If you uncheck this option, the bookmark does not appear in the Home section of the WebVPN portal.

Step 12 (Optional) Choose **Advanced Options** 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.

Adding a URL for a Predefined Application Template

This option simplifies bookmark creation with users selecting a predefined ASDM template that contains the pre-filled necessary values for certain well-defined applications.

Prerequisites

Predefined application templates are currently available for the following applications only:

- Citrix XenApp
- Citrix XenDesktop
- Domino WebAccess
- Microsoft Outlook Web Access 2010

- Microsoft Sharepoint 2007
- Microsoft SharePoint 2010

Detailed Steps

Step 1 Enter a name for the bookmark to display for the user.

Step 2 As a subtitle, provide additional user-visible text that describes the bookmark entry.

Step 3 Use the **Thumbnail** drop-down menu to select an icon to associate with the bookmark on the end-user portal.

Step 4 Click **Manage** to import or export images to use as thumbnails.

Step 5 (Optional) Select the **Place this bookmark on the VPN home page** check box.

Step 6 In the **Select Auto Sign-on Application** list, click the required application. The available applications are:

- Citrix XenApp
- Citrix XenDesktop
- Domino WebAccess
- Microsoft Outlook Web Access 2010
- Microsoft Sharepoint 2007
- Microsoft SharePoint 2010

Step 7 Enter the URL of the page which is loaded before the login page. This page will require user interaction to proceed to the login screen. The URL will allow * to substitute an arbitrary number of symbols, for example http*://www.example.com/test.

Step 8 Enter the **Pre-login Page Control ID**. This is the ID of the control / tag that will get a click event on the pre-login page URL to proceed to the login page.

Step 9 Enter the **Application Parameters**. Depending on the application these may include the following:

- **Protocol**. HTTP or HTTPs.
- **Host Name**. For example `www.cisco.com`.
- **Port Number**. The port used by the application.
- **URL Path Appendix**. For example `/Citrix/XenApp`. This is normally auto-populated.
- **Domain**. The domain to connect to
- **User Name**. The SSL VPN Variable to use as a user name. Click **Select Variable** to choose a different variable.
- P**assword**. The SSL VPN Variable to use as a password. Click **Select Variable** to choose a different variable.

Step 10 (Optional) Click **Preview** to view the template output. You can click **Edit** to modify the template.

Step 11 Click **OK** to make your changes. Alternatively, click **Cancel** to abandon your changes.

Adding a Bookmark for an Auto Sign-On Application

This option lets you create a bookmark for any complex auto-sign on application.

Prerequisites

Configuring auto sign-on applications requires two steps:

1. Define the bookmark with some basic initial data and without the POST parameters. Save and assign the bookmark to use in a group or user policy.
2. Edit the bookmark again. Use the capture function to capture the SSL VPN parameters and edit them in the bookmark.

Detailed Steps

Step 1 Enter a name for the bookmark to display for the user.

Step 2 Use the URL 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.

Step 3 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**

Step 4 As a subtitle, provide additional user-visible text that describes the bookmark entry.

Step 5 Use the **Thumbnail** drop-down menu to select an icon to associate with the bookmark on the end-user portal.

Step 6 Click **Manage** to import or export images to use as thumbnails.

Step 7 (Optional) Select the **Place this bookmark on the VPN home page** check box.

Step 8 Enter the **Login Page URL**. Wildcards can be used in the URL you enter. For example, you can enter `http*://www.example.com/myurl*`.

Step 9 Enter the **Landing Page URL**. The ASA requires the Landing Page to be configured to detect a successful login to the application.

Step 10 (Optional) Enter a **Post Script**. Some Web applications, such as Microsoft Outlook Web Access, may execute a JavaScript to change the request parameters before the log-on form is submitted. The **Post Script** field enables you to enter JavaScript for such applications.

Step 11 Add the required **Form Parameters**. For each required SSL VPN Variable, click **Add**, enter a **Name**, and select a variable from the list. You can click **Edit** to change parameters and **Delete** to remove them.

Step 12 Enter the URL of the page which is loaded before the login page. This page will require user interaction to proceed to the login screen. The URL will allow * to substitute an arbitrary number of symbols, for example http*://www.example.com/test.

Step 13 Enter the **Pre-login Page Control ID**. This is the ID of the control / tag that will get a click event on the pre-login page URL to proceed to the login page.

Step 14 Click **OK** to make your changes. Alternatively, click **Cancel** to abandon your changes.

When you edit the bookmark you can use the HTML Parameter Capture function to capture the VPN auto sign-on parameters. The bookmark must have been saved and assigned first to a group policy or user.

Enter the **SSL VPN Username** then click **Start Capture**. Then use a Web browser to start the VPN session and navigate to the intranet page. To complete the process, click Stop Capture. The parameters will then be available for editing and inserted in the bookmark.

Importing/Exporting 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.

Detailed Steps

Step 1 Identify the bookmark list by name. Maximum is 64 characters, no spaces.

Step 2 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 ASA.
- Remote server—Click to import a url list file that resides on a remote server accessible from the ASA.
- 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.

Importing/Exporting GUI Customization Objects (Web Contents)

This dialogue box lets you import and export web content objects. The names of the web content objects and their file types are displayed.

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.

Step 1 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 ASA.
- Remote server—Click to import a file that resides on a remote server accessible from the ASA.
- 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.

Step 2 Determine whether authentication is required to access the content.

The prefix to the path changes depending on whether you require authentication. The ASA uses /+CSCOE+/ for objects that require authentication, and /+CSCOU+/ for objects that do not. The ASA displays /+CSCOE+/ objects on the portal page only, while /+CSCOU+/ objects are visible and usable in either the logon or the portal pages.

Step 3 Click to import or export the file.

Adding/Editing Post Parameter

Use this pane to configure post parameters for bookmark entries and URL lists.

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.

Detailed Steps

Step 1 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:

Table 91-17 Clientless SSL VPN Variables

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.
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.

Table 91-17 Clientless SSL VPN Variables

No.	Variable Substitution	Definition
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 ASA 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 ASA 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 ASA 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 ASA 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 ASA 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. Without writing a script or uploading anything, an administrator can specify which homepage in the group policy to connect with via smart tunnel.

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.

Configuration Example for 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.

Configuration Example for 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 ASA 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

Configuration Example for Customizing External Ports

You can use the external portal feature to create your own portal instead of using the pre-configured one. If you set up your own portal, you can bypass the clientless portal and send a POST request to retrieve your portal.

Detailed Steps

Step 1 Choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Customization**. Highlight the desired customization and choose **Edit**.

Step 2 Check the Enable External Portal check box.

Step 3 In the URL field, enter the desired external portal so that POST requests are allowed.

Sending an Administrator's Alert to Clientless SSL VPN Users

To send an alert message to clientless SSL VPN users (for example, about connection status), 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**.

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

Configuration > Features > VPN > E-mail Proxy > AAA

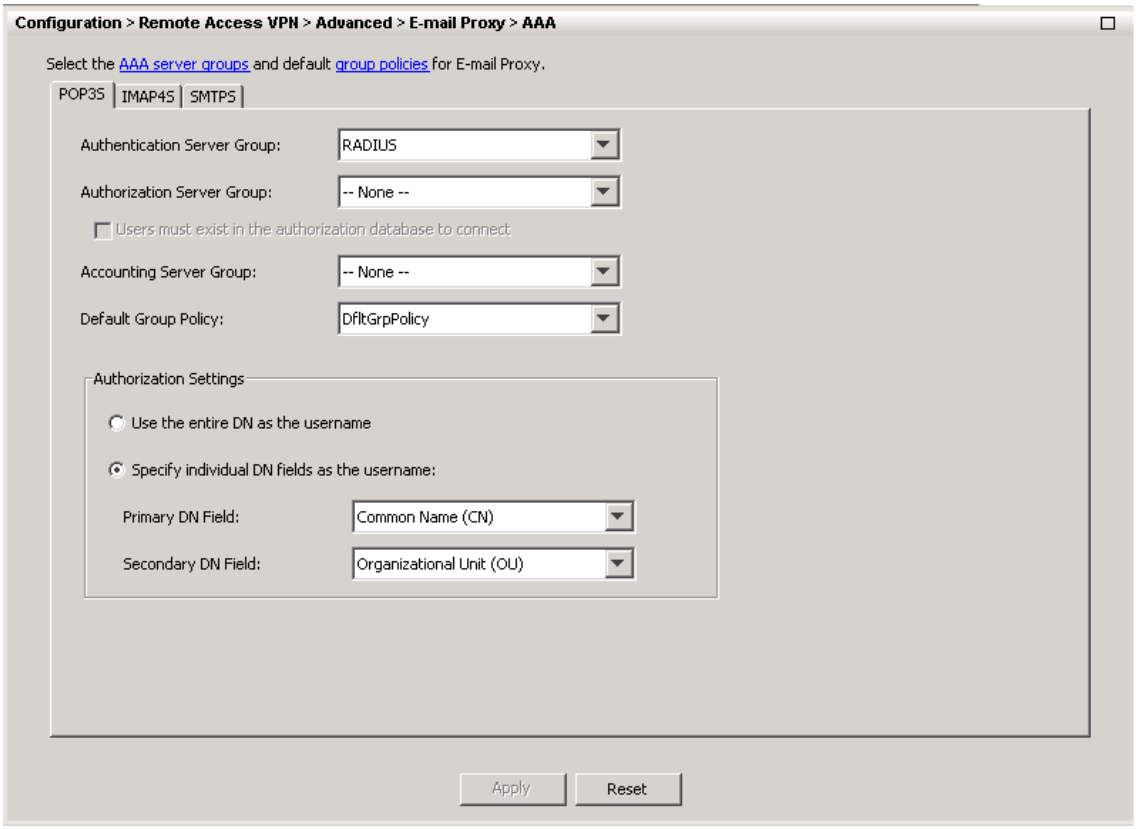

This panel has three tabs:

- POP3S Tab
- IMAP4S Tab
- SMTPS Tab

POP3S Tab

Configuration > Features > VPN > E-mail Proxy > AAA > POP3S Tab

The POP3S AAA panel associates AAA server groups and configures the default group policy for POP3S sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policies—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.

- Authentication Server Group—Select the authentication server group for POP3S user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for POP3S (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for POP3S user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for POP3S user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to POP3S users when AAA does not return a CLASSID attribute. The length must be between 4 and 15 alphanumeric characters. If you do not specify a default group policy, and there is no CLASSID, the ASA can not establish the session.
- Authorization Settings—Lets you set values for usernames that the ASA 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.

DN Field	Definition
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.

IMAP4S Tab

Configuration > Features > VPN > E-mail Proxy > AAA > IMAP4S Tab

The IMAP4S AAA panel associates AAA server groups and configures the default group policy for IMAP4S sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policy—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.
- Authentication Server Group—Select the authentication server group for IMAP4S user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for IMAP4S (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for IMAP4S user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for IMAP4S user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to IMAP4S users when AAA does not return a CLASSID attribute. If you do not specify a default group policy, and there is no CLASSID, the ASA can not establish the session.
- Authorization Settings—Lets you set values for usernames that the ASA recognizes for IMAP4S authorization. This applies to IMAP4S users that authenticate with digital certificates and require LDAP or RADIUS authorization.
  - User the entire DN as the username—Select to use the fully qualified domain name for IMAP4S authorization.
  - Specify individual DN fields as the username—Select to specify specific DN fields for user authorization.

    You can choose two DN fields, primary and secondary. For example, if you choose EA, users authenticate according to their e-mail address. Then a user with the Common Name (CN) John Doe and an e-mail address of johndoe@cisco.com cannot authenticate as John Doe or as johndoe. He must authenticate as johndoe@cisco.com. If you choose EA and O, John Does must authenticate as johndoe@cisco.com *and* Cisco. Systems, Inc.

– **Primary DN Field**—Select the primary DN field you want to configure for IMAP4S authorization. The default is CN. Options include the following:

DN Field	Definition
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

– Secondary DN Field—(Optional) Select the secondary DN field you want to configure for IMAP4S authorization. The default is OU. Options include all of those in the preceding table, with the addition of None, which you select if you do not want to include a secondary field.

SMTPS Tab

Configuration > Features > VPN > E-mail Proxy > AAA > SMTPS Tab

The SMTPS AAA panel associates AAA server groups and configures the default group policy for SMTPS sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policy—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.

- Authentication Server Group—Select the authentication server group for SMTPS user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for SMTPS (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for SMTPS user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for SMTPS user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to SMTPS users when AAA does not return a CLASSID attribute. If you do not specify a default group policy, and there is no CLASSID, the ASA can not establish the session.
- Authorization Settings—Lets you set values for usernames that the ASA 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.

DN Field	Definition
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.

Access

Configuration > VPN > E-Mail Proxy > Access

The E-mail Proxy Access screen lets you identify interfaces on which to configure e-mail proxy. You can configure 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.

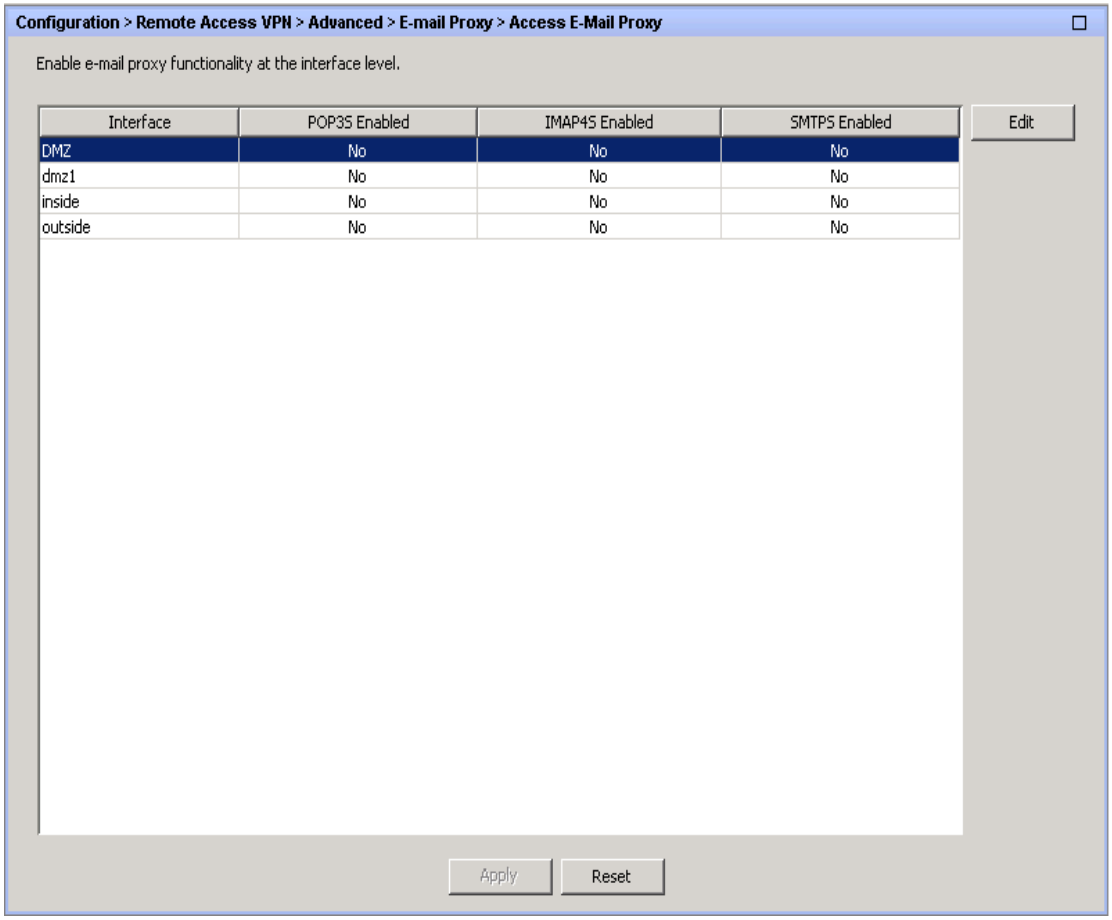

Fields

- 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.

Edit E-Mail Proxy Access

Configuration > VPN > E-Mail Proxy > Access > Edit E-Mail Proxy Access

The E-mail Proxy Access screen lets you identify interfaces on which to configure e-mail proxy. You can configure e-mail proxies on individual interfaces, and you can configure e-mail proxies for one interface and then apply your settings to all interfaces.

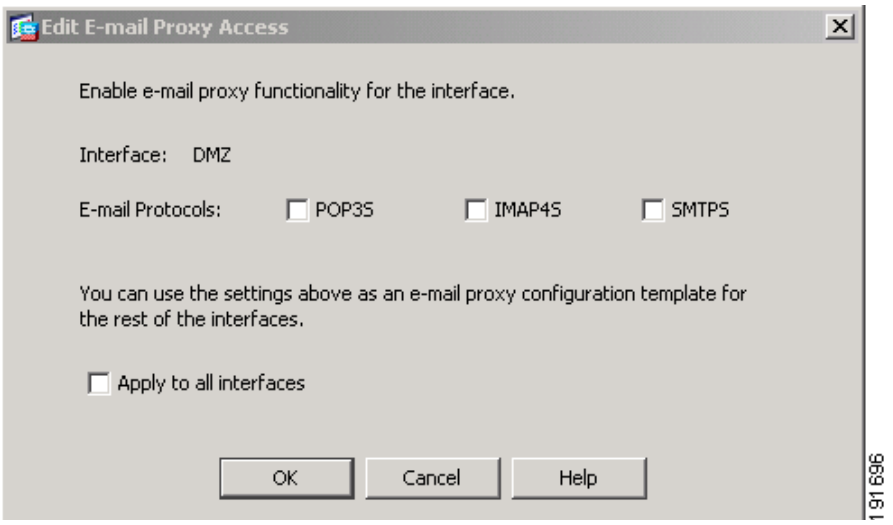

Fields

- Interface—Displays the name of the selected interface.
- POP3S Enabled—Select to enable POP3S for the interface.
- IMAP4S Enabled—elect to enable IMAP4S for the interface.
- SMTPS Enabled—Select to enable SMTPS for the interface.
- Apply to all interface—Select to apply the settings for the current interface to all configured interfaces.

Authentication

Configuration > Features > VPN >E-mail Proxy > Authentication

This panel lets you configure authentication methods for e-mail proxy sessions.

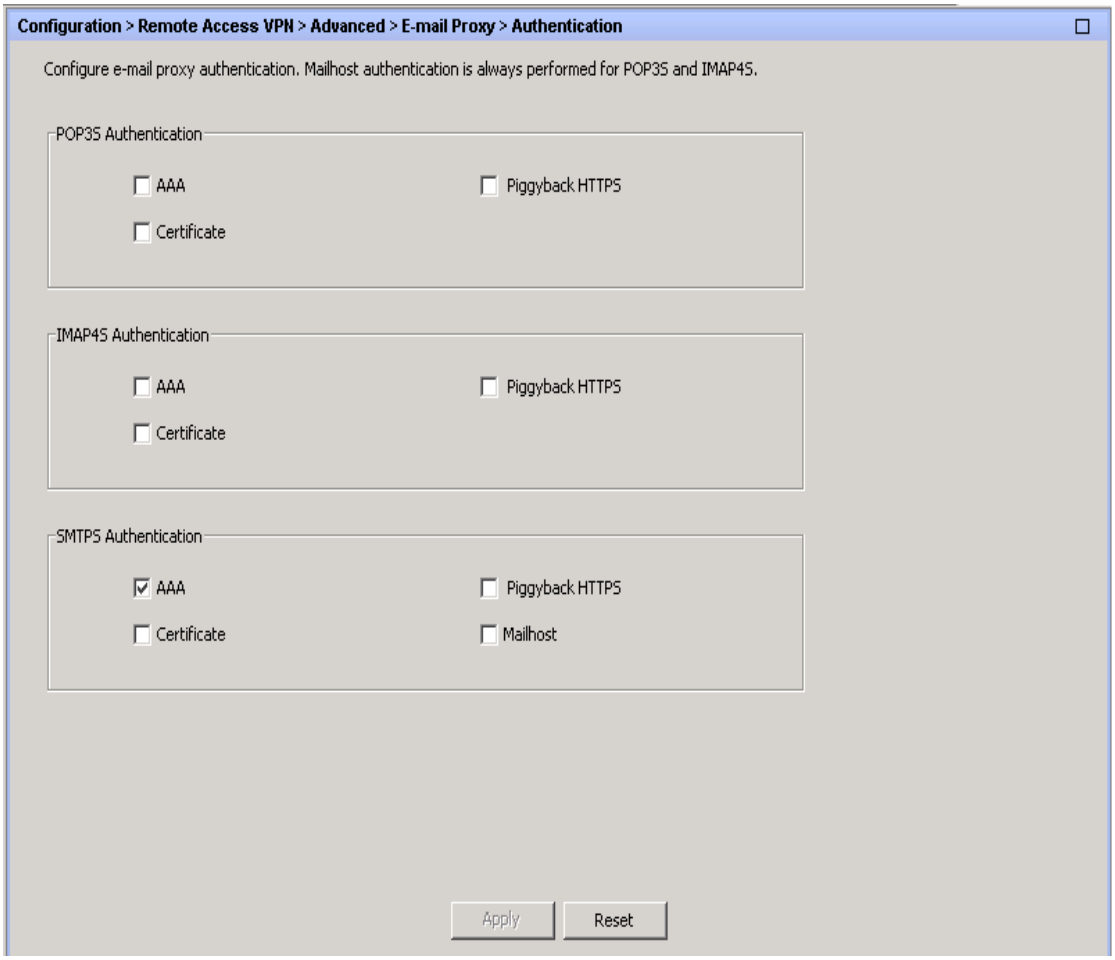

Fields

POP3S/IMAP4S/SMTPS Authentication—Let you configure authentication methods for each of the e-mail proxy types. You can select multiple methods of authentication.

- AAA—Select to require AAA authentication. This option requires a configured AAA server. The user presents a username, server and password. Users must present both the VPN username and the e-mail username, separated by the VPN Name Delimiter, only if the usernames are different from each other.
- Certificate—Certificate authentication does not work for e-mail proxies in the current ASA 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 ASA, 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.

Default Servers

Configuration > Features > VPN > E-mail Proxy > Default Servers

This panel lets you identify proxy servers to the ASA. Enter the IP address and port of the appropriate proxy server.

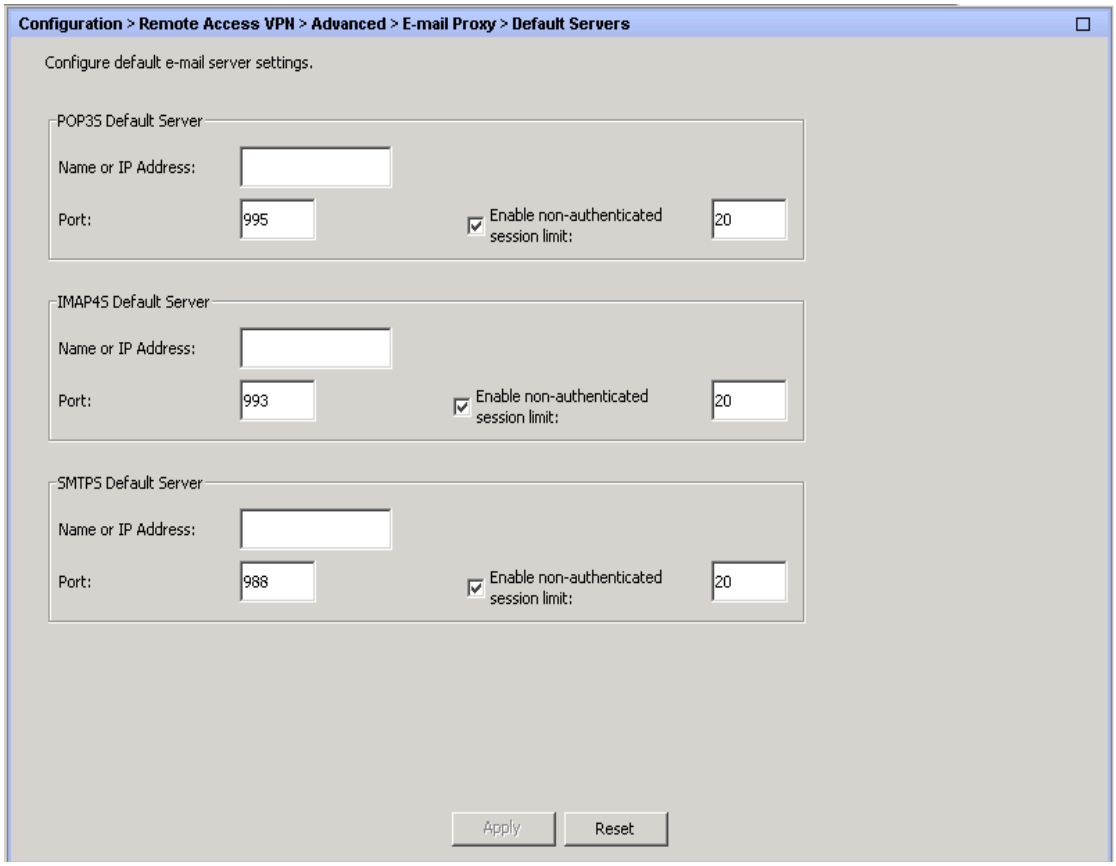

Fields

- 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 ASA 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 ASA 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 ASA 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.

Delimiters

Configuration > Features > VPN > E-mail Proxy > Delimiters

This panel lets you configure username/password delimiters and server delimiters for e-mail proxy authentication.

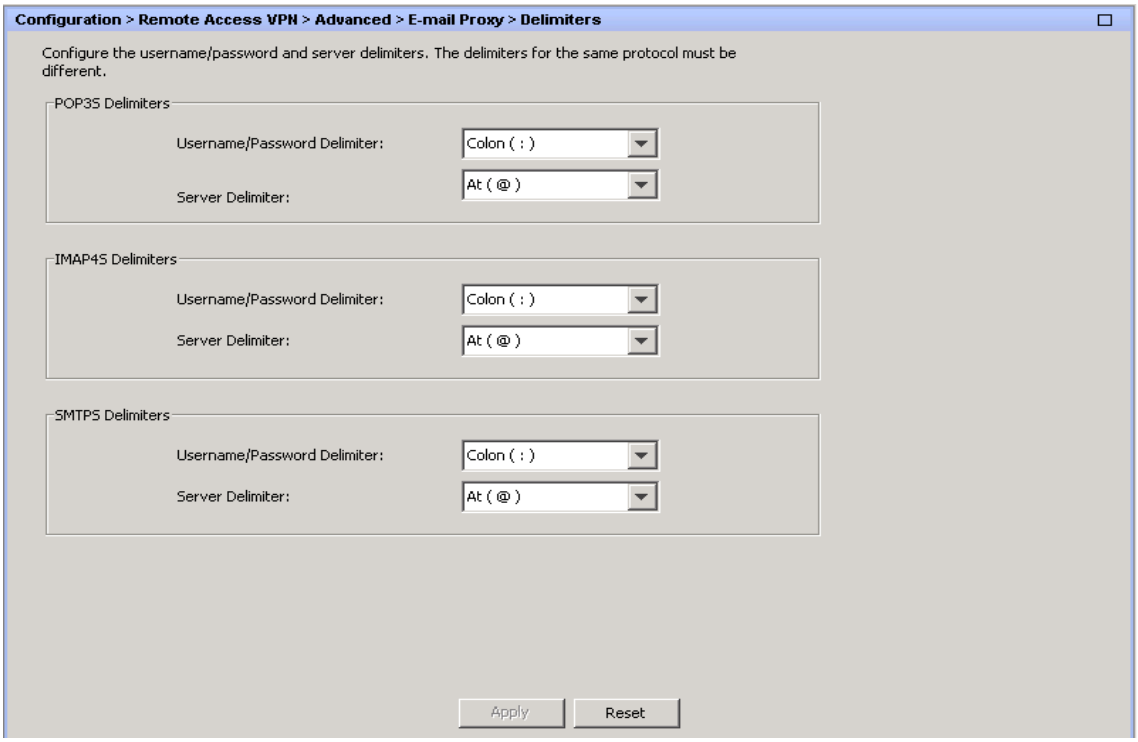

Fields

- POP3S/IMAP4S/SMTPS Delimiters—Let you configure username/password and server delimiters for each of the e-mail proxies.
  - Username/Password Delimiter—Select a delimiter to separate the VPN username from the e-mail username. Users need both usernames when using AAA authentication for e-mail proxy and the VPN username and e-mail username are different. Users enter both usernames, separated by the delimiter you configure here, and also the e-mail server name, when they log in to an e-mail proxy session.

Note Passwords for 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.

Monitoring VPN

This chapter describes how to use VPN monitoring parameters and statistics for the following:

- VPN statistics for specific Network (Client) Remote Access, Site-to-Site VPN, 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 ASA.

IPsec Tunnels

Monitoring > VPN > VPN Connection Graphs > 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.

Sessions

Monitoring > VPN > VPN Connection Graphs > 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 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, 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.

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 Window

Monitoring > VPN > VPN Statistics > Sessions

Use this pane to view VPN 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.
  - All Remote Access—Shows the number of remote access sessions.
  - Site-to-Site—Shows the number of LAN-to-LAN sessions.
  - Clientless SSL VPN—Shows the number of clientless browser-based VPN sessions.
  - AnyConnect 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 statistics 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 ASA 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. You can filter by any of these sessions: IPsec Site-to-Site, All Remote Access, AnyConnect Client, Clientless SSL VPN, IPsec (IKEv1) Remote Access, OSPFv3 IPsec, or Email Proxy.

    The column headings in the results table change depending on the Session type you choose.
  - 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.

    If you choose a Session type filter of IPsec Site-to-Site, AnyConnect Client, Clientless SSL VPN, or OSPFv3 IPsec then you would be able to filter by Assigned IP Address Type or Public IP Address Type.
  - Filter value—Enter the value that corresponds to the filter you are using.

    If you are filtering by Assigned IP Address Type or Public IP Address type, specify the IP address type in this field, either IPv4 or IPv6.
  - 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.

- All 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 hardware 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 ASA.

- IPsec Site-to-Site—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 ASA.
- 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 ASA.
- 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 ASA.
- 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 ASA.

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 ASA sends an ICMP Echo Request message to a selected host. If the host is reachable, it returns an Echo Reply message, and the ASA 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 ASA 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.

Viewing Active AnyConnect Sessions

Step 1 Select **Monitoring > VPN > VPN Statistics > Sessions**.

Step 2 In the Filter By field, select **AnyConnect Client**.

Step 3 In the Session Filter field (unlabeled), next to the Filter By field, select the session type you want to use to further refine your filter. Then, enter a value in the Session Value field (unlabeled) to the right of the Session Filter field. These are the available session filters and session values:

Session Filter	Session Value
Username	Then enter the username you want to sort by.
Assigned IP Address	The assigned IP address is the IP address assigned by the ASA to the AnyConnect Client connection. Enter the IPv4 or IPv6 address you want to sort by.
Assigned IP Address Type	Choose IP version 4 or IP version 6.
Public Address	The public IP address is the IP address assigned by your enterprise to the endpoint. Enter the IPv4 or IPv6 address you want to sort by.
Public Address Type	Choose IP version 4 or IP version 6.
Encryption	Select the encryption type for the session.
Connection Status	Select Active or Inactive.

Step 4 Cick **Filter**.

Viewing VPN Sessions Details

Monitoring > VPN > VPN Statistics > 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 ASA.
- NAC Result and Posture Token—The ASDM displays values in this column only if you configured Network Admission Control on the ASA.

  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 ASA.
  - Non-Responsive—The remote host did not respond to the EAPoUDP Hello message.
  - Hold-off—The ASA 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 ASA 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 pane displays the following columns:

- ID—Unique ID dynamically assigned to the session. The ID serves as the ASA 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 ASA 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 ASA 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 ASA. The Redirect URL is an optional part of the access policy payload. The ASA 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 ASA 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.

Cluster Loads

Monitoring > VPN > VPN Statistics > 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 ASA 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.

Crypto Statistics

Monitoring > VPN > VPN Statistics > Crypto Statistics

This pane displays the crypto statistics for currently active user and administrator sessions on the ASA. 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.

Compression Statistics

Monitoring > VPN > VPN Statistics > Compression Statistics

This pane displays the compression statistics for currently active user and administrator sessions on the ASA. 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.

Encryption Statistics

Monitoring > VPN > VPN Statistics > Encryption Statistics

This pane shows the data encryption algorithms used by currently active user and administrator sessions on the ASA. 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 ASA was last booted or reset.
- Refresh—Updates the statistics shown in the Encryption Statistics table.

Global IKE/IPsec Statistics

Monitoring > VPN > VPN Statistics > Global IKE/IPSec Statistics

This pane displays the global IKE/IPsec statistics for currently active user and administrator sessions on the ASA. 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.

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 ASA.
- 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 ASA configuration supports clientless hosts, the Access Control Server downloads the access policy associated with clientless hosts to the ASA for these peers. Otherwise, the ASA assigns the NAC default policy.
- Hold-off—Number of peers for which the ASA 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 ASA. 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 ASA, 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.

Protocol Statistics

Monitoring > VPN > VPN Statistics > Protocol Statistics

This pane displays the protocols used by currently active user and administrator sessions on the ASA. 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 ASA was last booted or reset.
- Refresh—Updates the statistics shown in the Protocol Statistics table.

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 ASA forwards all traffic to the specified VLAN.

Field

- Active VLAN Mapping Sessions—Number of VPN sessions assigned to an egress VLAN.

SSO Statistics for Clientless SSL VPN Session

Monitoring > VPN > WebVPN > SSO Statistics

This pane displays the single sign-on statistics for currently active SSO servers configured for the ASA.

Note These statistics are for SSO with SiteMinder and SAML Browser Post Profile servers only.

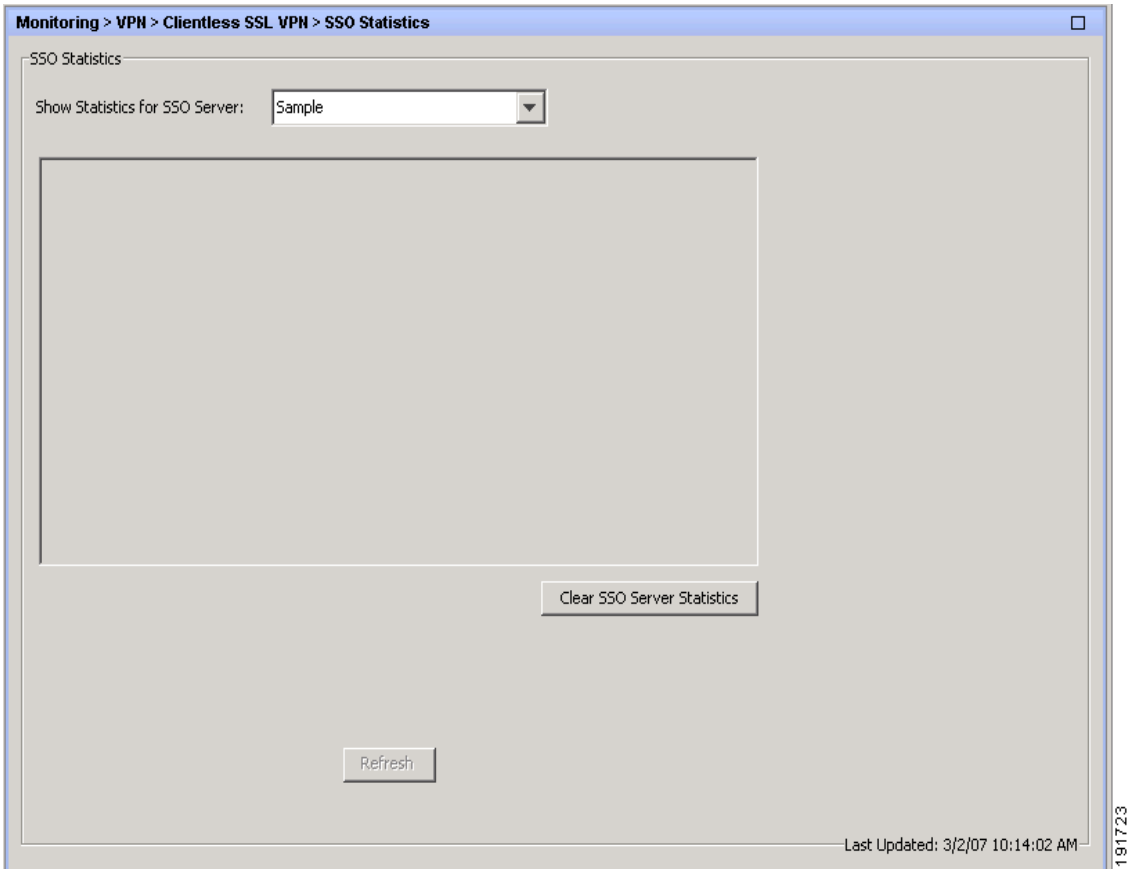

Fields

- 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.

VPN Connection Status for the Easy VPN Client

Use this pane to view the status of the ASA configured as an Easy VPN client. This features applies to the ASA 5505 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.

Configuring SSL Settings

SSL

Configuration > Properties > SSL

The ASA 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 ASA uses to negotiate as a server. You can make only one selection.

Options for Server SSL versions include the following:

Any	The ASA accepts SSL version 2 client hellos, and negotiates either SSL version 3 or TLS version 1.
Negotiate SSL V3	The ASA accepts SSL version 2 client hellos, and negotiates to SSL version 3.
Negotiate TLS V1	The ASA 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 ASA uses to negotiate as a client. You can make only one selection.

Options for Client SSL versions include the following:

any	The ASA 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 ASA 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 ASA 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 ASA 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.

Edit SSL Certificate

Configuration > Properties > SSL > Edit SSL Trustpoint

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.

SSL Certificates

In this pane, you can require that device management sessions require user certificates for SSL authentication.

Configuration > Properties > SSL > Edit SSL Trustpoint

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.

PART 16

Configuring Logging, SNMP, and Smart Call Home

Configuring Logging

This chapter describes how to configure and manage logs for the ASA and ASASM and includes the following sections:

- Information About Logging, page 96-1
- Licensing Requirements for Logging, page 96-5
- Prerequisites for Logging, page 96-6
- Guidelines and Limitations, page 96-6
- Configuring Logging, page 96-6
- Monitoring the Logs, page 96-24
- Feature History for Logging, page 96-28

Information About Logging

System logging is a method of collecting messages from devices to a server running a syslog daemon. Logging to a central syslog server helps in aggregation of logs and alerts. Cisco devices can send their log messages to a UNIX-style syslog service. A syslog service accepts messages and stores them in files, or prints them according to a simple configuration file. This form of logging provides protected long-term storage for logs. Logs are useful both in routine troubleshooting and in incident handling.

The ASA system logs provide you with information for monitoring and troubleshooting the ASA. With the logging feature, you can do the following:

- Specify which syslog messages should be logged.
- Disable or change the severity level of a syslog message.
- Specify one or more locations where syslog messages should be sent, including an internal buffer, one or more syslog servers, ASDM, an SNMP management station, specified e-mail addresses, or to Telnet and SSH sessions.
- Configure and manage syslog messages in groups, such as by severity level or class of message.
- Specify whether or not a rate-limit is applied to syslog generation.
- Specify what happens to the contents of the internal log buffer when it becomes full: overwrite the buffer, send the buffer contents to an FTP server, or save the contents to internal flash memory.
- Filter syslog messages by locations, severity level, class, or a custom message list.

This section includes the following topics:

- Logging in Multiple Context Mode, page 96-2
- Analyzing Syslog Messages, page 96-2
- Syslog Message Format, page 96-3
- Severity Levels, page 96-3
- Message Classes and Range of Syslog IDs, page 96-4
- Filtering Syslog Messages, page 96-4
- Sorting in the Log Viewers, page 96-4
- Using Custom Message Lists, page 96-5
- Using Clustering, page 96-5

Logging in Multiple Context Mode

Each security context includes its own logging configuration and generates its own messages. If you log in to the system or admin context, and then change to another context, messages you view in your session are only those messages 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 ASA and ASASM 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.

Analyzing Syslog Messages

The following are some examples of the type of information you can obtain from a review of various syslog messages:

- Connections that are allowed by ASA and ASASM security policies. These messages help you spot holes that remain open in your security policies.
- Connections that are denied by ASA and ASASM security policies. These messages show what types of activity are being directed toward your secured inside network.
- Using the ACE deny rate logging feature shows attacks that are occurring on your ASA or ASA Services Module.
- IDS activity messages can show attacks that have occurred.
- User authentication and command usage provide an audit trail of security policy changes.
- Bandwidth usage messages show each connection that was built and torn down as well as the duration and traffic volume used.
- Protocol usage messages show the protocols and port numbers used for each connection.
- Address translation audit trail messages record NAT or PAT connections being built or torn down, which are useful if you receive a report of malicious activity coming from inside your network to the outside world.

Syslog Message Format

Syslog messages begin with a percent sign (%) and are structured as follows:

```
%ASA Level Message_number: Message_text
```

Field descriptions are as follows:

ASA	The syslog message facility code for messages that are generated by the ASA and ASASM. This value is always ASA.
Level	1 through 7. The level reflects the severity of the condition described by the syslog message—the lower the number, the more severe the condition. See Table 96-1 for more information.
Message_number	A unique six-digit number that identifies the syslog message.
Message_text	A text string that describes the condition. This portion of the syslog message sometimes includes IP addresses, port numbers, or usernames.

Severity Levels

Table 96-1 lists the syslog message severity levels. You can assign custom colors to each of the severity levels to make it easier to distinguish them in the ASDM log viewers. To configure syslog message color settings, either choose the **Tools > Preferences > Syslog** tab or, in the log viewer itself, click **Color Settings** on the toolbar.

Table 96-1 Syslog Message Severity Levels

Level Number	Severity Level	Description
0	**emergencies**	System is unusable.
1	**alert**	Immediate action is needed.
2	**critical**	Critical conditions.
3	**error**	Error conditions.
4	**warning**	Warning conditions.
5	**notification**	Normal but significant conditions.
6	**informational**	Informational messages only.
7	**debugging**	Debugging messages only.

Note The ASA and ASASM do not generate syslog messages with a severity level of zero (emergencies). This level is provided in the **logging** command for compatibility with the UNIX syslog feature but is not used by the ASA.

Message Classes and Range of Syslog IDs

For a list of syslog message classes and the ranges of syslog message IDs that are associated with each class, see the syslog messages guide.

Filtering Syslog Messages

You can filter generated syslog messages so that only certain syslog messages are sent to a particular output destination. For example, you could configure the ASA and ASASM to send all syslog messages to one output destination and to send a subset of those syslog messages to a different output destination.

Specifically, you can configure the ASA and ASASM so that syslog messages are directed to an output destination according to the following criteria:

- Syslog message ID number
- Syslog message severity level
- Syslog message class (equivalent to a functional area of the ASA and ASASM)

You customize these criteria by creating a message list that you can specify when you set the output destination. Alternatively, you can configure the ASA or ASASM to send a particular message class to each type of output destination independently of the message list.

You can use syslog message classes in two ways:

- Specify an output location for an entire category of syslog messages using the **logging class** command.
- Create a message list that specifies the message class using the **logging list** command.

The syslog message class provides a method of categorizing syslog messages by type, equivalent to a feature or function of the ASA and ASASM. For example, the vpnc class denotes the VPN client.

All syslog messages in a particular class share the same initial three digits in their syslog message ID numbers. For example, all syslog message IDs that begin with the digits 611 are associated with the vpnc (VPN client) class. Syslog messages associated with the VPN client feature range from 611101 to 611323.

In addition, most of the ISAKMP syslog messages have a common set of prepended objects to help identify the tunnel. These objects precede the descriptive text of a syslog message when available. If the object is not known at the time that the syslog message is generated, the specific *heading = value* combination does not appear.

The objects are prefixed as follows:

Group = *groupname*, Username = *user*, IP = *IP_address*

Where the group is the tunnel-group, the username is the username from the local database or AAA server, and the IP address is the public IP address of the remote access client or L2L peer.

Sorting in the Log Viewers

You can sort messages in all ASDM log viewers (that is, the Real-Time Log Viewer, the Log Buffer Viewer, and the Latest ASDM Syslog Events Viewer). To sort tables by multiple columns, click the header of the first column that you want to sort by, then press and hold down the **Ctrl** key and at the same

time, click the headers of the other column(s) that you want to include in the sort order. To sort messages chronologically, select both the date and time columns; otherwise, the messages are sorted only by date (regardless of the time) or only by time (regardless of the date).

When you sort messages in the Real-Time Log Viewer and in the Latest ASDM Syslog Events Viewer, the new messages that come in appear in the sorted order, instead of at the top, as they normally would be. That is, they are mixed in with the rest of the messages.

Using Custom Message Lists

Creating a custom message list is a flexible way to exercise control over which syslog messages are sent to which output destination. In a custom syslog message list, you specify groups of syslog messages using any or all of the following criteria: severity level, message IDs, ranges of syslog message IDs, or message class.

For example, you can use message lists to do the following:

- Select syslog messages with the severity levels of 1 and 2 and send them to one or more e-mail addresses.
- Select all syslog messages associated with a message class (such as ha) and save them to the internal buffer.

A message list can include multiple criteria for selecting messages. However, you must add each message selection criterion with a new command entry. It is possible to create a message list that includes overlapping message selection criteria. If two criteria in a message list select the same message, the message is logged only once.

Using Clustering

Syslog messages are an invaluable tool for accounting, monitoring, and troubleshooting in a clustering environment. Each ASA unit in the cluster (up to eight units are allowed) generates syslog messages independently; certain **logging** commands then enable you to control header fields, which include a timestamp and device ID. The syslog server uses the device ID to identify the syslog generator. You can use the **logging device-id** command to generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster.

Note To monitor syslog messages from units in a cluster, you must open an ASDM session to each of the units that you want to monitor.

Licensing Requirements for Logging

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites for Logging

Logging has the following prerequisites:

- The syslog server must run a server program called syslogd. Windows (except for Windows 95 and Windows 98) provides a syslog server as part of its operating system. For Windows 95 and Windows 98, you must obtain a syslogd server from another vendor.
- To view logs generated by the ASA or ASASM, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the ASA and ASASM generate messages but does not save them to a location from which you can view them. You must specify each different logging output destination separately. For example, to designate more than one syslog server as an output destination, specify separate entries in the Syslog Server pane for each syslog server.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context modes.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

- Sending syslogs over TCP is not supported on a standby ASA.
- The ASA supports the configuration of 16 syslog servers with the **logging host** command in single context mode. In multiple context mode, the limitation is 4 servers per context.

Configuring Logging

This section describes how to configure logging and includes the following topics:

- Enabling Logging, page 96-7
- Configuring an Output Destination, page 96-7

Note The minimum configuration depends on what you want to do and what your requirements are for handling syslog messages in the ASA and ASASM.

Enabling Logging

To enable logging, perform the following steps:

Path	Purpose
Choose one of the following: • **Home > Latest ASDM Syslog Messages > Enable Logging** • **Configuration > Device Management > Logging > Logging Setup** • **Monitoring > Real-Time Log Viewer > Enable Logging** • **Monitoring > Log Buffer > Enable Logging**	Enables logging.

What to Do Next

See the "Configuring an Output Destination" section on page 96-7.

Configuring an Output Destination

To optimize syslog message usage for troubleshooting and performance monitoring, we recommend that you specify one or more locations where syslog messages should be sent, including an internal log buffer, one or more external syslog servers, ASDM, an SNMP management station, the console port, specified e-mail addresses, or Telnet and SSH sessions.

This section includes the following topics:

- Sending Syslog Messages to an External Syslog Server, page 96-8
- Configuring FTP Settings, page 96-9
- Configuring Logging Flash Usage, page 96-9
- Configuring Syslog Messaging, page 96-10
- Editing Syslog ID Settings, page 96-10
- Including a Device ID in Non-EMBLEM Formatted Syslog Messages, page 96-11
- Sending Syslog Messages to the Internal Log Buffer, page 96-12
- Saving an Internal Log Buffer to Flash, page 96-13
- Viewing and Copying Logged Entries with the ASDM Java Console, page 96-13
- Sending Syslog Messages to an E-mail Address, page 96-13
- Adding or Editing E-Mail Recipients, page 96-14
- Configuring the Remote SMTP Server, page 96-14
- Viewing Syslog Messages in ASDM, page 96-15
- Applying Message Filters to a Logging Destination, page 96-15
- Applying Logging Filters, page 96-15

- Adding or Editing a Message Class and Severity Filter, page 96-16
- Adding or Editing a Syslog Message ID Filter, page 96-17
- Sending Syslog Messages to the Console Port, page 96-17
- Sending Syslog Messages to a Telnet or SSH Session, page 96-17
- Creating a Custom Event List, page 96-17
- Generating Syslog Messages in EMBLEM Format to a Syslog Server, page 96-18
- Adding or Editing Syslog Server Settings, page 96-19
- Generating Syslog Messages in EMBLEM Format to Other Output Destinations, page 96-19
- Changing the Amount of Internal Flash Memory Available for Logs, page 96-20
- Configuring the Logging Queue, page 96-20
- Sending All Syslog Messages in a Class to a Specified Output Destination, page 96-21
- Enabling Secure Logging, page 96-21
- Including the Device ID in Non-EMBLEM Format Syslog Messages, page 96-21
- Including the Date and Time in Syslog Messages, page 96-22
- Disabling a Syslog Message, page 96-22
- Changing the Severity Level of a Syslog Message, page 96-22
- Limiting the Rate of Syslog Message Generation, page 96-22
- Assigning or Changing Rate Limits for Individual Syslog Messages, page 96-23
- Adding or Editing the Rate Limit for a Syslog Message, page 96-23
- Editing the Rate Limit for a Syslog Severity Level, page 96-24

Sending Syslog Messages to an External Syslog Server

You can archive messages according to the available disk space on the external syslog server, and manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

To send syslog messages to an external syslog server, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Logging Setup**.

Step 2 Check the **Enable logging** check box to turn on logging for the main ASA.

Step 3 Check the **Enable logging on the failover standby unit** check box to turn on logging for the standby ASA, if available.

Step 4 Check the **Send debug messages as syslogs** check box to redirect 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 5 Check the **Send syslogs in EMBLEM format** check box to enable EMBLEM format so that it is used for all logging destinations, except syslog servers.

Step 6 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 are 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 7 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 8 Click **Configure FTP Settings** to identify the FTP server and configure the FTP parameters used to save the buffer content. For more information, see the "Configuring FTP Settings" section on page 96-9.

Step 9 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 10 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 the "Configuring Logging Flash Usage" section on page 96-9.

Note This option is only available in single routed or transparent mode.

Step 11 In the Queue Size field, specify the queue size for system logs that are to be viewed in the ASA or ASASM.

Configuring FTP Settings

To specify the configuration for the FTP server that is used to save the log buffer content, 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**.

Configuring Logging Flash Usage

To specify the limits for saving the log buffer content to internal flash memory, 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.

Configuring Syslog Messaging

To configure syslog messaging, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Syslog Setup**.

Step 2 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 3 To add the date and time in each syslog message sent, check the **Include timestamp in syslogs** check box.

Step 4 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 5 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 6 To configure syslog messages to include a device ID, click **Advanced**. For more information, see the "Editing Syslog ID Settings" section on page 96-10 and the "Including a Device ID in Non-EMBLEM Formatted Syslog Messages" section on page 96-11.

Editing Syslog ID Settings

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)

  Note Using a severity level of zero is not recommended.

- 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.

Including a Device ID in Non-EMBLEM Formatted Syslog Messages

To include a device ID in non-EMBLEM formatted syslog messages, 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 of the ASA
- Interface IP address

  Choose the interface name that corresponds to the selected IP address from the drop-down list.

  If you are using clustering, check the **In an ASA cluster, always use master's IP address for the selected interface** check box.

- String

  In the User-Defined ID field, specify an alphanumeric, user-defined string.

- ASA cluster name

Step 3 Click **OK** to close this dialog box.

Sending Syslog Messages to the Internal Log Buffer

To send syslog messages to the internal log buffer, perform the following steps:

	Path	Purpose
Step 1	Choose one of the following: • **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters** • **Configuration > Device Management > Logging > Logging Filters**	Specifies which syslog messages should be sent to the internal log buffer, which serves as a temporary storage location. New messages are appended to the end of the list. When the buffer is full, that is, when the buffer wraps, old messages are overwritten as new messages are generated, unless you configure the ASA and ASASM to save the full buffer to another location. To empty the internal log buffer, choose **Monitoring > Logging > Log Buffer > View**. Then in the Log Buffer pane, choose **File > Clear Internal Log Buffer**.
Step 2	**Configuration > Device Management > Logging > Logging Setup**	Changes the size of the internal log buffer. The buffer size is 4 KB.
Step 3	Choose one of the following options:	
	Configuration > Device Management > Logging > Logging Setup > Configure Flash Usage	When saving the buffer content to another location, the ASA and ASASM create log files with names that use the following time-stamp format: `LOG-YYYY-MM-DD-HHMMSS.TXT` where *YYYY* is the year, *MM* is the month, *DD* is the day of the month, and *HHMMSS* is the time in hours, minutes, and seconds. The ASA and ASASM continue to save new messages to the internal log buffer and saves the full log buffer content to the internal flash memory.
	Configuration > Device Management > Logging > Logging Setup > Configure FTP Settings	When saving the buffer content to another location, the ASA and ASASM create log files with names that use the following time-stamp format: `LOG-YYYY-MM-DD-HHMMSS.TXT` where *YYYY* is the year, *MM* is the month, *DD* is the day of the month, and *HHMMSS* is the time in hours, minutes, and seconds. The ASA and ASASM continue saving new messages to the internal log buffer and saves the full log buffer content to an FTP server.
	Configuration > Device Management > Logging > Logging Setup > Configure FTP Settings	Identifies the FTP server on which you want to store log buffer content.
	Configuration > Device Management > Logging > Logging Setup > Configure FTP Settings	Saves the current log buffer content to the internal flash memory.

Saving an Internal Log Buffer to Flash

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**.

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.

Step 5 To continue, see the "Monitoring Performance" section on page 73-8.

Sending Syslog Messages to an E-mail Address

To send syslog messages to an e-mail address, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > E-Mail Setup**.

Step 2 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 3 Click **Add** to enter a new e-mail address recipient of the specified syslog messages. For more information, see the "Adding or Editing E-Mail Recipients" section on page 96-14.

Step 4 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 the "Applying Logging Filters" section on page 96-15.

Step 5 Click **Edit** to modify an existing severity level of the syslog messages that are sent to this recipient. For more information, see the "Adding or Editing E-Mail Recipients" section on page 96-14.

Step 6 Click **OK** to close this dialog box.

Step 7 To continue, see the "Configuring the Remote SMTP Server" section on page 96-14.

Adding or Editing E-Mail Recipients

To add or edit e-mail recipients and severity levels, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > E-mail Setup**.

Step 2 Click **Add** or **Edit** to display the Add/Edit E-Mail Recipient dialog box.

Step 3 Enter the destination e-mail address, and choose the syslog severity level from the drop-down list. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)

Note Using a severity level of zero is not recommended.

- 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)

Note The severity level used to filter messages for the destination e-mail address is the higher of the severity level specified in the Add/Edit E-Mail Recipient dialog box and the global filter set for all e-mail recipients in the Logging Filters pane.

Step 4 Click **OK** to close this dialog box.

The added or revised entry appears in the E-mail Recipients pane.

Step 5 Click **Apply** to save your changes to the running configuration.

Configuring the Remote SMTP Server

To configure the remote SMTP server to which e-mail alerts and notifications are sent in response to specific events, perform the following steps:

Step 1 Choose **Configuration > Device Setup > Logging > SMTP**.

Step 2 Enter the IP address of the primary SMTP server.

Step 3 (Optional) Enter the IP address of the standby SMTP server, then click **Apply** to save your changes to the running configuration.

Viewing Syslog Messages in ASDM

To view the latest syslog messages that have been sent to ASDM, choose **Home > Latest ASDM Syslog Messages**. The ASA or ASASM sets aside a buffer area for syslog messages waiting to be sent to ASDM and saves messages in the buffer as they occur. The ASDM log buffer is a different buffer than the internal log buffer. When the ASDM log buffer is full, the ASA or ASASM deletes the oldest syslog message to make room in the buffer for new ones. Deleting the oldest syslog message to make room for new ones is the default setting in ASDM.

Applying Message Filters to a Logging Destination

To apply message filters to a logging destination, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Logging Filters**.

Step 2 Choose the name of the logging destination to which you want to apply a filter. Available logging destinations are as follows:

- ASDM
- Console port
- E-Mail
- Internal buffer
- SNMP server
- Syslog server
- Telnet or SSH session

  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 logging destination, or whether logging is disabled for all event classes. The third column lists the event class to use to filter messages for that logging destination. For more information, see the "Adding or Editing a Message Class and Severity Filter" section on page 96-16 and the "Adding or Editing a Syslog Message ID Filter" section on page 96-17.

Step 3 Click **Edit** to display the Edit Logging Filters dialog box. To apply, edit, or disable filters, see the "Applying Logging Filters" section on page 96-15.

Applying Logging Filters

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 the "Creating a Custom Event List" section on page 96-17.

Step 5 Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

Step 6 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note Using a severity level of zero is not recommended.

- 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**.

The selected logging destination for a filter appears at the top.

Adding or Editing a Message Class and Severity Filter

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 change according to the device mode that you are using.

Step 2 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note Using a severity level of zero is not recommended.

- 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.

Adding or Editing a Syslog Message ID Filter

To add or edit a syslog message ID filter, see the "Editing Syslog ID Settings" section on page 96-10.

Sending Syslog Messages to the Console Port

To send syslog messages to the console port, enter the following command:

Path	Purpose
Choose one of the following: • **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters** • **Configuration > Device Management > Logging > Logging Filters**	Specifies which syslog messages should be sent to the console port.

Sending Syslog Messages to a Telnet or SSH Session

To send syslog messages to a Telnet or SSH session, perform the following steps:

	Path	Purpose
Step 1	Choose one of the following: • **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters** • **Configuration > Device Management > Logging > Logging Filters**	Specifies which syslog messages should be sent to a Telnet or SSH session.
Step 2	**Configuration > Device Management > Logging > Logging Setup**	Enables logging to the current session only.

Creating a Custom Event List

You use the following three criteria to define an event list:

- Event Class
- Severity
- Message ID

To create a custom list of events to send to a specific logging destination (for example, an SNMP server), perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Event Lists**.

Step 2 Click **Add** to display the Add Event List dialog box.

Step 3 In the Name field, enter the name of the event list. No spaces are allowed.

Step 4 In the Event Class/Severity area, click **Add** to display the Add Class and Severity Filter dialog box.

Step 5 Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

Step 6 Choose the severity level from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note Using a severity level of zero is not recommended.

- 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 **OK** to close this dialog box.

Step 8 In the Message ID Filters area, click **Add** to display the Add Syslog Message ID Filter dialog box.

Step 9 In the Message IDs field, enter a syslog message ID or range of IDs (for example, 101001-199012) to include in the filter.

Step 10 Click **OK** to close this dialog box.

The event of interest appears in the list. To change this entry, click **Edit**.

Generating Syslog Messages in EMBLEM Format to a Syslog Server

To generate syslog messages in EMBLEM format to a syslog server, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Syslog Server**.

Step 2 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.

Note You can set up a maximum of four syslog servers per security context (up to a total of 16).

Step 3 Specify the number of messages that are allowed to be queued on the ASA or ASASM when a syslog server is busy. A zero value means an unlimited number of messages may be queued.

Step 4 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 5 To continue, see the "Adding or Editing Syslog Server Settings" section on page 96-19.

Adding or Editing Syslog Server Settings

To add or edit syslog server settings, 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 ASA or ASASM.

Step 4 Enter the port number used by the syslog server to communicate with the ASA or ASASM.

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.

Generating Syslog Messages in EMBLEM Format to Other Output Destinations

To generate syslog messages in EMBLEM format to other output destinations, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Logging Setup	Sends syslog messages in EMBLEM format to output destinations other than a syslog server, such as Telnet or SSH sessions.

Changing the Amount of Internal Flash Memory Available for Logs

To change the amount of internal flash memory available for logs, perform the following steps:

	Path	Purpose
Step 1	**Configuration > Device Management > Logging > Logging Setup**	Specifies the maximum amount of internal flash memory available for saving log files. By default, the ASA can use up to 1 MB of internal flash memory for log data. The minimum amount of internal flash memory that must be free for the ASA and ASASM to save log data is 3 MB. If a log file being saved to internal flash memory would cause the amount of free internal flash memory to fall below the configured minimum limit, the ASA or ASASM deletes the oldest log files to ensure that the minimum amount of memory remains free after saving the new log file. If there are no files to delete or if, after all old files have been deleted, free memory is still below the limit, the ASA or ASASM fails to save the new log file.
Step 2	**Configuration > Device Management > Logging > Logging Setup**	Specifies the minimum amount of internal flash memory that must be free for the ASA or ASASM to save a log file.

Configuring the Logging Queue

To configure the logging queue, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Logging Setup	Specifies the number of syslog messages that the ASA and ASASM can hold in its queue before sending them to the configured output destination. The ASA and ASASM have a fixed number of blocks in memory that can be allocated for buffering syslog messages while they are waiting to be sent to the configured output destination. The number of blocks required depends on the length of the syslog message queue and the number of syslog servers specified. The default queue size is 512 syslog messages. The queue size is limited only by block memory availability. Valid values are from 0 to 8192 messages, depending on the platform. If the logging queue is set to zero, the queue is the maximum configurable size (8192 messages), depending on the platform. The maximum queue size by platform is as follows: • ASA-5505—1024 • ASA-5510—2048 • On all other platforms—8192

Sending All Syslog Messages in a Class to a Specified Output Destination

To send all syslog messages in a class to a specified output destination, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Logging Filters	Overrides the configuration in the specified output destination command. For example, if you specify that messages at severity level 7 should go to the internal log buffer and that ha class messages at severity level 3 should go to the internal log buffer, then the latter configuration takes precedence. To specify that a class should go to more than one destination, choose a different filtering option for each output destination.

Enabling Secure Logging

To enable secure logging, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Syslog Server	Enables secure logging.

Including the Device ID in Non-EMBLEM Format Syslog Messages

To include the device ID in non-EMBLEM format syslog messages, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration	Configures the ASA or ASASMto include a device ID in non-EMBLEM-format syslog messages. You can specify only one type of device ID for syslog messages. If you enable the logging device ID for the admin context in multiple context mode, 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. **Note** In an ASA cluster, always use the master unit IP address for the selected interface. **Note** If enabled, the device ID does not appear in EMBLEM-formatted syslog messages nor in SNMP traps.

Including the Date and Time in Syslog Messages

To include the date and time in syslog messages, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Syslog Setup	Specifies that syslog messages should include the date and time that they were generated.

Disabling a Syslog Message

To disable a specified syslog message, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Syslog Setup	Prevents the ASA or ASASMfrom generating a particular syslog message.

Changing the Severity Level of a Syslog Message

To change the severity level of a syslog message, enter the following command:

Path	Purpose
Configuration > Device Management > Logging > Syslog Setup	Specifies the severity level of a syslog message.

Limiting the Rate of Syslog Message Generation

To limit the rate of syslog message generation, perform the following steps:

Step 1 Choose **Configuration > Device Management > Logging > Rate Limit**.

Step 2 Choose the logging level (message severity level) to which you want to assign rate limits. Severity levels are defined as follows:

Description	Severity Level
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

Description	Severity Level
Informational	6—Informational messages only
Debugging	7—Debugging messages only

Step 3 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.

Step 4 To continue, see the "Assigning or Changing Rate Limits for Individual Syslog Messages" section on page 96-23.

Assigning or Changing Rate Limits for Individual Syslog Messages

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.

Step 2 To continue, see the "Adding or Editing the Rate Limit for a Syslog Message" section on page 96-23.

Step 3 To change the rate limit of a specific syslog message, click **Edit** to display the Edit Rate Limit for Syslog Message dialog box.

Step 4 To continue, see the "Editing the Rate Limit for a Syslog Severity Level" section on page 96-24.

Adding or Editing the Rate Limit for a 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.

Editing the Rate Limit for a Syslog Severity Level

To change the rate limit of a specified syslog severity level, perform the following steps:

Step 1 Enter the maximum number of messages at this severity 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 severity level, and click **OK**.

The selected message severity level appears.

Note To allow an unlimited number of messages, leave both the Number of Messages and Time Interval fields blank.

Monitoring the Logs

This section includes the following topics:

- Filtering Syslog Messages Through the Log Viewers, page 96-25
- Editing Filtering Settings, page 96-27
- Executing Certain Commands Using the Log Viewers, page 96-27

To monitor the logs in the log buffer or in real-time and assist in monitoring the system performance, choose one of the following:

Path	Purpose
Choose one of the following: • **Monitoring > Logging > Log Buffer > View** • **Monitoring > Logging > Real-Time Log Viewer > View**	Shows syslog messages, including the severity level. **Note** The maximum number of syslog messages that are available to view is 1000, which is the default setting. The maximum number of syslog messages that are available to view is 2000. Displays the message explanations, additional details, and recommended actions to take, if necessary, to resolve an error in a separate window. Provides text search within messages and message filtering. Allows creation of a reverse access control rule that performs the opposite action of the access control rule that originally generated the syslog message. Reverse access control rules can be applied *only* to syslog messages 106100, 106023, 338001 through 338004, 338201, and 338202. Provides the ability to use the following options in the log viewers from the Tools menu: Ping, Traceroute, Whois, and DNS Lookup. Provides sorting of messages in each column shown. Allows detailed message filtering based on the syslog ID, date and time, severity level, source and destination IP addresses, source and destination ports, and description listed. Displays popup help in the Build Filter dialog box.

Filtering Syslog Messages Through the Log Viewers

You can filter syslog messages based on one or multiple values that correspond to any column of the Real-Time Log Viewer and the Log Buffer Viewer.

To filter syslog messages through one of the log viewers, perform the following steps:

Step 1 Choose one of the following:

- **Monitoring > Logging > Real-Time Log Viewer > View**
- **Monitoring > Logging > Log Buffer > View**

Step 2 In either the Real-Time Log Viewer or the Log Buffer Viewer dialog box, click **Build Filter** on the toolbar.

Step 3 In the Build Filter dialog box, specify the filtering criteria to apply to syslog messages:

a. In the Date and Time area, choose one of the following three options: real-time, a specific time, or a time range. If you chose a specific time, indicate the time by entering the number and choosing hours or minutes from the drop-down list. If you chose a time range, in the Start Time field, click the drop-down arrow to display a calendar. Choose a start date and a start time from the drop-down list, then click **OK**. In the End Time field, click the drop-down arrow to display a calendar. Choose an end date and an end time from the drop-down list, then click **OK**.

b. Enter a valid severity level in the Severity field. Alternatively, click the **Edit** icon on the right of the Severity field. In the Severity dialog box, click the severity levels in the list on which you want to filter. To include severity levels 1-7, click **All**. Click **OK** to display these settings in the Build Filter dialog box. For additional information about the correct input format to use, click the **Info** icon on the right of the Severity field.

c. Enter a valid syslog ID in the Syslog ID field. Alternatively, click the **Edit** icon on the right of the Syslog ID field. In the Syslog ID dialog box, choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the Build Filter dialog box. Click **Delete** to remove these settings and enter new ones. For additional information about the correct input format to use, click the **Info** icon on the right of the Syslog ID field.

d. Enter a valid source IP address in the Source IP Address field, or click the **Edit** icon on the right of the Source IP Address field. In the Source IP Address dialog box, choose a single IP address or a specified range of IP addresses, then click **Add**. To exclude a specific IP address or range of IP addresses, check the **Do not include (exclude) this address or range** check box. Click **OK** to display these settings in the Build Filter dialog box. Click **Delete** to remove these settings and enter new ones. For additional information about the correct input format to use, click the **Info** icon on the right of the Source IP Address field.

e. Enter a valid source port in the Source Port field, or click the **Edit** icon on the right of the Source Port field. In the Source Port dialog box, choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the Build Filter dialog box. Click **Delete** to remove these settings and enter new ones. For additional information about the correct input format to use, click the **Info** icon on the right of the Source Port field.

f. Enter a valid destination IP address in the Destination IP Address field, or click the **Edit** icon on the right of the Destination IP Address field. In the Destination IP Address dialog box, choose a single IP address or a specified range of IP addresses, then click **Add**. To exclude a specific IP address or range of IP addresses, check the **Do not include (exclude) this address or range** check box. Click **OK** to display these settings in the Build Filter dialog box. Click **Delete** to remove these settings and enter new ones. For additional information about the correct input format to use, click the **Info** icon on the right of the Destination IP Address field.

g. Enter a valid destination port in the Destination Port field, or click the **Edit** icon on the right of the Destination Port field. In the Destination Port dialog box, choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the Build Filter dialog box. Click **Delete** to remove these settings and enter new ones. For additional information about the correct input format to use, click the **Info** icon on the right of the Destination Port field.

h. Enter filtering text for the Description field. The text may be any string of one or more characters, including a regular expression. However, semicolons are not valid characters, and this setting is case-sensitive. Multiple entries must be separated by commas.

i. Click **OK** to add the filter settings you have just specified to the Filter By drop-down list in the log viewers. The filter strings follow a specific format. The prefix FILTER: designates all custom filters that appear in the Filter By drop-down list. You may still type random text into this field.

   The following table shows examples of the format used.

Build Filter Example	Filter String Format
Source IP = 192.168.1.1 or 0.0.0.0 Source Port = 67	FILTER: srcIP=192.168.1.1,0.0.0.0;srcPort=67;
Severity = Informational Destination IP = 1.1.1.1 through 1.1.1.10	FILTER: sev=6;dstIP=1.1.1.1-1.1.1.10;

Syslog ID not in the range 725001 through 725003	FILTER: sysID=!725001-725003;
Source IP = 1.1.1.1 Description = Built outbound	FILTER: srcIP=1.1.1.1;descr=Built outbound

Step 4 To filter syslog messages, choose one of the settings in the Filter By drop-down list, then click **Filter** on the toolbar. This setting also applies to all future syslog messages. To clear all filters, click **Show All** on the toolbar.

Note You cannot save filters that you have specified with the Build Filter dialog box. These filters are valid only for the ASDM session during which they were created.

Editing Filtering Settings

To edit filtering settings that you created using the Build Filter dialog box, perform the following steps:

Choose one of the following:

- Revise a filter directly by entering the changes in the Filter By drop-down list.
- Choose a filter in the Filter By drop-down list, then click **Build Filter** to display the Build Filter dialog box. To remove the current filter settings and enter new ones, click **Clear Filter**. Otherwise, change the settings that appear, and click **OK**.

  Note These filter settings apply only to those defined in the Build Filter dialog box.

- To stop filtering and show all syslog messages, click **Show All** on the toolbar.

Executing Certain Commands Using the Log Viewers

You can execute the following commands using either of the log viewers: **ping**, **traceroute**, **whois**, and **dns lookup**.

To execute any of these commands, perform the following steps:

Step 1 Choose one of the following:

- **Monitoring > Logging > Real-Time Log Viewer > View**
- **Monitoring Logging > Log Buffer > View**

Step 2 From the Real-Time Log Viewer or Log Buffer pane, click **Tools**, then choose the command that you want to execute. Alternatively, you can right-click a specific syslog message that is listed to display a context menu, then choose the command that you want to execute.

The Entering command dialog box appears, with the command that you selected automatically showing in the drop-down list.

Step 3 Enter either the source or destination IP address of the selected syslog message in the Address field, then click **Go**.

The command output appears in the area provided.

Step 4 Click **Clear** to remove the output, and choose another command to execute from the drop-down list. Repeat Step 3, if necessary. Click **Close** when you are done.

Feature History for Logging

Table 96-2 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 96-2 Feature History for Logging

Feature Name	Platform Releases	Feature Information
Logging	7.0(1)	Provides ASA network logging information through various output destinations, and includes the option to view and save log files. We introduced the following screen: Configuration > Device Management > Logging > Logging Setup.
Rate limit	7.0(4)	Limits the rate at which syslog messages are generated. We modified the following screen: Configuration > Device Management > Logging > Rate Limit.
Logging list	7.2(1)	Creates a logging list to use in other commands to specify messages by various criteria (logging level, event class, and message IDs). We modified the following screen: Configuration > Device Management > Logging > Event Lists.
Secure logging	8.0(2)	Specifies that the connection to the remote logging host should use SSL/TLS. This option is valid only if the protocol selected is TCP. We modified the following screen: Configuration > Device Management > Logging > Syslog Server.
Logging class	8.0(4), 8.1(1)	Added support for the ipaa event class of logging messages. We modified the following screen: Configuration > Device Management > Logging > Logging Filters.
Logging class and saved logging buffers	8.2(1)	Added support for the dap event class of logging messages. Added support to clear the saved logging buffers (ASDM, internal, FTP, and flash). We modified the following screen: Configuration > Device Management > Logging > Logging Setup.
Password encryption	8.3(1)	Added support for password encryption.
Log viewers	8.3(1)	The source and destination IP addresses were added to the log viewers.

Table 96-2 Feature History for Logging (continued)

Feature Name	Platform Releases	Feature Information
Enhanced logging and connection blocking	8.3(2)	When you configure a syslog server to use TCP, and the syslog server is unavailable, the ASA blocks new connections that generate syslog messages until the server becomes available again (for example, VPN, firewall, and cut-through-proxy connections). This feature has been enhanced to also block new connections when the logging queue on the ASA is full; connections resume when the logging queue is cleared. This feature was added for compliance with Common Criteria EAL4+. Unless required, we recommended allowing connections when syslog messages cannot be sent or received. To allow connections, continue to check the **Allow user traffic to pass when TCP syslog server is down** check box on the Configuration > Device Management > Logging > Syslog Servers pane. We introduced the following syslog messages: 414005, 414006, 414007, and 414008. We did not modify any ASDM screens.
Syslog message filtering and sorting	8.4(1)	Support has been added for the following: • Syslog message filtering based on multiple text strings that correspond to various columns • Creation of custom filters • Column sorting of messages. For detailed information, see the ASDM configuration guide. We modified the following screens: Monitoring > Logging > Real-Time Log Viewer > View. Monitoring > Logging > Log Buffer Viewer > View. This feature interoperates with all ASA versions.
Clustering	9.0(1)	Added support for syslog message generation in a clustering environment on the ASA 5580 and 5585-X. We modified the following screen: Configuration > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration.

Configuring NetFlow Secure Event Logging (NSEL)

This chapter describes how to configure NSEL, a security logging mechanism that is built on NetFlow Version 9 technology, and how to handle events and syslog messages through NSEL.

This chapter includes the following sections:

- Information About NSEL, page 97-1
- Licensing Requirements for NSEL, page 97-4
- Prerequisites for NSEL, page 97-4
- Guidelines and Limitations, page 97-4
- Configuring NSEL, page 97-5
- Monitoring NSEL, page 97-7
- Where to Go Next, page 97-7
- Additional References, page 97-7
- Feature History for NSEL, page 97-8

Information About NSEL

This section includes the following topics:

- Using NSEL and Syslog Messages, page 97-2
- Using NSEL in Clustering, page 97-3

The ASA and ASASM support NetFlow Version 9 services. For more information about NetFlow services, see the "RFCs" section on page 97-8.

The ASA and ASASM implementations of NSEL provide a stateful, IP flow tracking method that exports only those records that indicate significant events in a flow. In stateful flow tracking, tracked flows go through a series of state changes. NSEL events are used to export data about flow status and are triggered by the event that caused the state change.

The significant events that are tracked include flow-create, flow-teardown, flow-denied (excluding those flows that are denied by EtherType ACLs), and flow-update. In addition, the ASA and ASASM implementation of NSEL generates periodic NSEL events and flow-update events to provide periodic byte counters over the duration of the flow. These events are usually time-driven, which makes them more in line with traditional Netflow; however, these events may also be triggered by state changes in the flow.

Note The flow-update event feature is available only in Version 8.4(5). It is not available in Version 9.0(1) or later.

Each NSEL record has an event ID and an extended event ID field, which describes the flow event.

The ASA and ASASM implementations of NSEL provide the following major functions:

- Tracks flow-create, flow-teardown, and flow-denied events, and generates appropriate NSEL data records.
- Defines and exports templates that describe the progression of a flow. Templates describe the format of the data records that are exported through NetFlow. Each event has several record formats or templates associated with it.
- Tracks configured NSEL collectors and delivers templates and data records to these configured NSEL collectors through NetFlow over UDP only.
- Sends template information periodically to NSEL collectors. Collectors receive template definitions, normally before receiving flow records.
- Filters NSEL events based on the traffic and event type through Modular Policy Framework, then sends records to different collectors. Traffic is matched based on the order in which classes are configured. After a match is found, no other classes are checked. The supported event types are flow-create, flow-denied, flow-teardown, and all. Records can be sent to different collectors. For example, with two collectors, you can do the following:
  - Log all flow-denied events that match access list 1 to collector 1.
  - Log all flow-create events to collector 1.
  - Log all flow-teardown events to collector 2.
- Delays the export of flow-create events.

Using NSEL and Syslog Messages

Table 97-1 lists the syslog messages that have an equivalent NSEL event, event ID, and extended event ID. The extended event ID provides more detail about the event (for example, which ACL—ingress or egress—has denied a flow).

Note Enabling NetFlow to export flow information makes the syslog messages that are listed in Table 97-1 redundant. In the interest of performance, we recommend that you disable redundant syslog messages, because the same information is exported through NetFlow.

Table 97-1 Syslog Messages and Equivalent NSEL Events

Syslog Message	Description	NSEL Event ID	NSEL Extended Event ID
106100	Generated whenever an ACL is encountered.	1—Flow was created (if the ACL allowed the flow). 3—Flow was denied (if the ACL denied the flow).	0—If the ACL allowed the flow. 1001—Flow was denied by the ingress ACL. 1002—Flow was denied by the egress ACL.
106015	A TCP flow was denied because the first packet was not a SYN packet.	3—Flow was denied.	1004—Flow was denied because the first packet was not a TCP SYN packet.
106023	When a flow was denied by an ACL attached to an interface through the **access-group** command.	3—Flow was denied.	1001—Flow was denied by the ingress ACL. 1002—Flow was denied by the egress ACL.
302013, 302015, 302017, 302020	TCP, UDP, GRE, and ICMP connection creation.	1—Flow was created.	0—Ignore.
302014, 302016, 302018, 302021	TCP, UDP, GRE, and ICMP connection teardown.	2—Flow was deleted.	0—Ignore. > 2000—Flow was torn down.
313001	An ICMP packet to the device was denied.	3—Flow was denied.	1003—To-the-box flow was denied because of configuration.
313008	An ICMP v6 packet to the device was denied.	3—Flow was denied.	1003—To-the-box flow was denied because of configuration.
710003	An attempt to connect to the device interface was denied.	3—Flow was denied.	1003—To-the-box flow was denied because of configuration.

Note When NSEL and syslog messages are both enabled, there is no guarantee of chronological ordering between the two logging types.

Using NSEL in Clustering

Each ASA establishes its own connection to the collector(s) using its unit local IP address. The fields in the header of the export packet include the system up time, UNIX time (synchronized across the cluster), and sequence number. These fields are all local to an individual ASA. The NSEL collector uses the combination of the source port of the packet to separate different exporters.

Each ASA manages and advertises its template independently. Because the ASA supports in-cluster upgrades, different units may run different image versions at a certain point in time. As a result, the template that each ASA supports may be different.

Note Clustering is available on the ASA 5580 and 5585-X only. For more information about clustering, see Chapter 10, “Configuring a Cluster of ASAs.”

Licensing Requirements for NSEL

Model	License Requirement
All models	Base License.

Prerequisites for NSEL

NSEL has the following prerequisites:

- IP address and hostname assignments must be unique throughout the NetFlow configuration.
- You must have at least one configured collector before you can use NSEL.
- You must configure NSEL collectors before you can configure filters via Modular Policy Framework.

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 **class-map**, **match access-list**, and **match any** commands.

Additional Guidelines and Limitations

- If you have previously configured flow-export actions using the **flow-export enable** command, and you upgrade to a later version, then your configuration is automatically converted to the new Modular Policy Framework **flow-export event-type** command, which is described under the **policy-map** command.
- Flow-export actions are not supported in interface-based policies. You can configure flow-export actions in a class-map only with the **match access-list**, **match any**, or **class-default** commands. You can only apply flow-export actions in a global service policy.
- To view bandwidth usage for NetFlow records (not available in real-time), you must use the threat detection feature.
- Only the ASA 5580 and 5585-X support clustering.
- NetFlow records are sent only for connections whose source and destination addresses are both IPv4 or both IPv6. NetFlow records are not sent for connections with mixed traffic, such as NAT64 or NAT46.

Configuring NSEL

This section describes how to configure NSEL and includes the following topics:

- Using NetFlow, page 97-5
- Matching NetFlow Events to Configured Collectors, page 97-6

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, 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 NetFlow Collector dialog box, and perform the following steps:

a. Choose the interface to which NetFlow packets will be sent from the drop-down list.

b. Enter the IP address or hostname and the UDP port number in the associated fields.

c. Click **OK**.

Step 4 To configure more collectors, repeat Step 2 for each additional collector.

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 Click **Apply** to save your changes. Click **Reset** to enter new settings.

What to Do Next

See the "Matching NetFlow Events to Configured Collectors" section on page 97-6.

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

a. Click **Add** to display the Add Service Policy Rule Wizard. For more information about service policy rules, see the "Adding a Service Policy Rule for Through Traffic" section on page 44-8.

b. Click the **Global - applies to all interfaces** radio button to apply the rule to the global policy. Click **Next**.

c. Check the **Source and Destination IP Address (uses ACL)** check box or the **Any traffic** check box as traffic match criteria, or click the **Use class-default as traffic class** radio button. Click **Next** to continue to the Rule Actions screen.

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 In the Rule Actions screen, click the **NetFlow** tab.

Step 4 To specify flow events, click **Add** to display the Add Flow Event dialog box, then perform the following steps:

a. Choose the flow event type from the drop-down list. Available events are created, torn down, denied, updated, or all.

Note The flow-update event is available in Version 8.4(5) only. It is not available in Version 9.0(1) or later.

b. Choose collectors to which you want events sent by checking the corresponding check boxes in the Send column.

c. To add, edit or delete collectors, or to configure other NetFlow settings (for example, syslog messages), click **Manage** to display the Manage NetFlow Collectors dialog box. Click **OK** to close the Manage NetFlow Collectors dialog box and return to the Add Flow Event dialog box. For more information about configuring collectors, see Step 3 of the "Using NetFlow" section on page 97-5.

Step 5 Click **OK** to close the Add Flow Event dialog box and return to the NetFlow tab.

Step 6 To change flow event entries, select an entry from the list, and click **Edit**. To remove flow event entries, select an entry from the list, and click **Delete**.

Step 7 Click **Finish** to exit the wizard.

Step 8 To edit a NetFlow service policy rule, perform the following steps:

a. Select it in the Service Policy Rules table, and click **Edit**.

b. Click the **Rule Actions** tab, then click the **NetFlow** tab.

What to Do Next

See the "Monitoring NSEL" section on page 97-7.

Monitoring NSEL

You can use syslog messages to help troubleshoot errors or monitor system usage and performance.You can view real-time syslog messages that have been saved in the log buffer in a separate window, which include an explanation of the message, details about the message, and recommended actions to take, if necessary, to resolve an error. For more information, see the "Using NSEL and Syslog Messages" section on page 97-2.

To monitor NSEL, see the following pane:

Path	Purpose
Tools > Command Line Interface Enter the **show flow-export counters** command, then click **Send**.	Shows runtime counters, including statistical data and error data, for NSEL.
Tools > Command Line Interface Enter the **show running-config flow-export** command, then click **Send**.	Shows the currently configured NetFlow commands.
Tools > Command Line Interface Enter the **show logging flow-export-syslogs** command, then click **Send**.	Lists all syslog messages that are captured by NSEL events.
Tools > Command Line Interface Enter the **show running-config logging** command, then click **Send**.	Shows disabled syslog messages, which are redundant syslog messages, because they export the same information through NetFlow.

Where to Go Next

To configure the syslog server, see Chapter 96, "Configuring Logging."

Additional References

For additional information related to implementing NSEL, see the following sections:

- Related Documents, page 97-8
- RFCs, page 97-8

Related Documents

Related Topic	Document Title
Using NSEL and Syslog Messages, page 97-2	*syslog messages guide*
Information about the implementation of NSEL on the ASA and ASA Services Module	*Cisco ASA 5500 Series Implementation Note for NetFlow Collectors* See the following article at https://supportforums.cisco.com/docs/DOC-6113.
Configuring NetFlow on the ASA and ASA Services Module using ASDM	See the following article at https://supportforums.cisco.com/docs/DOC-6114.

RFCs

RFC	Title
3954	Cisco Systems NetFlow Services Export Version 9

Feature History for NSEL

Table 97-2 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 97-2 Feature History for NSEL

Feature Name	Platform Releases	Feature Information
NetFlow	8.1(1)	The NetFlow feature enhances the ASA logging capabilities by logging flow-based events through the NetFlow protocol. NetFlow Version 9 services are used to export information about the progression of a flow from start to finish. The NetFlow implementation exports records that indicate significant events in the life of a flow. This implementation is different from traditional NetFlow, which exports data about flows at regular intervals. The NetFlow module also exports records about flows that are denied by access lists. You can configure an ASA 5580 to send the following events using NetFlow: flow create, flow teardown, and flow denied (only flows denied by ACLs are reported). We introduced the following screen: Configuration > Device Management > Logging > NetFlow.
NetFlow Filtering	8.1(2)	You can filter NetFlow events based on traffic and event type, then send records to different collectors. For example, you can log all flow-create events to one collector, and log flow-denied events to a different collector. For short-lived flows, NetFlow collectors benefit from processing a single event instead of two events: flow create and flow teardown. You can configure a delay before sending the flow-create event. If the flow is torn down before the timer expires, only the flow teardown event is sent. The teardown event includes all information regarding the flow; no loss of information occurs. We modified the following screen: Configuration > Firewall > Service Policy Rules.

Table 97-2 Feature History for NSEL (continued)

Feature Name	Platform Releases	Feature Information
NSEL	8.2(1)	The NetFlow feature has been ported to all available models of ASAs.
Clustering	9.0(1)	The NetFlow feature supports clustering.
NSEL		A new NetFlow error counter, source port allocation failure, has been added. Flow-update events have been introduced to provide periodic byte counters for flow traffic. You can change the time interval at which flow-update events are sent to the NetFlow collector. You can filter to which collectors flow-update records will be sent. We modified the following screens: Configuration > Firewall > Service Policy Rules > Add Service Policy Rule Wizard - Rule Actions > NetFlow > Add Flow Event Configuration > Device Management > Logging > NetFlow **Note** The flow-update event feature is available in Version 8.4(5) only; it is not available in Version 9.0(1) or later.

Configuring SNMP

This chapter describes how to configure SNMP to monitor the ASA and ASASM and includes the following sections:

- Information About SNMP, page 98-1
- Licensing Requirements for SNMP, page 98-4
- Prerequisites for SNMP, page 98-4
- Guidelines and Limitations, page 98-4
- Configuring SNMP, page 98-6
- Monitoring SNMP, page 98-9
- Where to Go Next, page 98-11
- Additional References, page 98-11
- Feature History for SNMP, page 98-13

Information About SNMP

SNMP is an application-layer protocol that facilitates the exchange of management information between network devices and is part of the TCP/IP protocol suite. This section describes SNMP and includes the following topics:

- Information About SNMP Terminology, page 98-2
- SNMP Version 3, page 98-2

The ASA and ASASM provide support for network monitoring using SNMP Versions 1, 2c, and 3, and supports the use of all three versions simultaneously. The SNMP agent running on the ASA interface lets you monitor the ASA and ASASM through network management systems (NMSs), such as HP OpenView. The ASA and ASASM support SNMP read-only access through issuance of a GET request. SNMP write access is not allowed, so you cannot make changes with SNMP. In addition, the SNMP SET request is not supported.

You can configure the ASA and ASASM to send traps, which are unsolicited messages from the managed device to the management station for certain events (event notifications) to an NMS, or you can use the NMS to browse the MIBs on the ASA. MIBs are a collection of definitions, and the ASA and ASASM maintain a database of values for each definition. Browsing a MIB means issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the NMS to determine values.

The ASA and ASASM have an SNMP agent that notifies designated management stations if events occur that are predefined to require a notification, for example, when a link in the network goes up or down. The notification it sends includes an SNMP OID, which identifies itself to the management stations. The ASA or ASASM SNMP agent also replies when a management station asks for information.

Information About SNMP Terminology

Table 98-1 lists the terms that are commonly used when working with SNMP:

Table 98-1 SNMP Terminology

Term	Description
Agent	The SNMP server running on the ASA. The SNMP agent has the following features: • Responds to requests for information and actions from the network management station. • Controls access to its Management Information Base, the collection of objects that the SNMP manager can view or change. • Does not allow set operations.
Browsing	Monitoring the health of a device from the network management station by polling required information from the SNMP agent on the device. This activity may include issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the network management station to determine values.
Management Information Bases (MIBs)	Standardized data structures for collecting information about packets, connections, buffers, failovers, and so on. MIBs are defined by the product, protocols, and hardware standards used by most network devices. SNMP network management stations can browse MIBs and request specific data or events be sent as they occur.
Network management stations (NMSs)	The PCs or workstations set up to monitor SNMP events and manage devices, such as the ASA and ASASM.
Object identifier (OID)	The system that identifies a device to its NMS and indicates to users the source of information monitored and displayed.
Trap	Predefined events that generate a message from the SNMP agent to the NMS. Events include alarm conditions such as linkup, linkdown, coldstart, warmstart, authentication, or syslog messages.

SNMP Version 3

This section describes SNMP Version 3 and includes the following topics:

- SNMP Version 3 Overview, page 98-3
- Security Models, page 98-3
- SNMP Groups, page 98-3
- SNMP Users, page 98-3
- SNMP Hosts, page 98-3
- Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software, page 98-4

SNMP Version 3 Overview

SNMP Version 3 provides security enhancements that are not available in SNMP Version 1 or SNMP Version 2c. SNMP Versions 1 and 2c transmit data between the SNMP server and SNMP agent in clear text. SNMP Version 3 adds authentication and privacy options to secure protocol operations. In addition, this version controls access to the SNMP agent and MIB objects through the User-based Security Model (USM) and View-based Access Control Model (VACM). The ASA and ASASM also support the creation of SNMP groups and users, as well as hosts, which is required to enable transport authentication and encryption for secure SNMP communications.

Security Models

For configuration purposes, the authentication and privacy options are grouped together into security models. Security models apply to users and groups, which are divided into the following three types:

- NoAuthPriv—No Authentication and No Privacy, which means that no security is applied to messages.
- AuthNoPriv—Authentication but No Privacy, which means that messages are authenticated.
- AuthPriv—Authentication and Privacy, which means that messages are authenticated and encrypted.

SNMP Groups

An SNMP group is an access control policy to which users can be added. Each SNMP group is configured with a security model, and is associated with an SNMP view. A user within an SNMP group must match the security model of the SNMP group. These parameters specify what type of authentication and privacy a user within an SNMP group uses. Each SNMP group name and security model pair must be unique.

SNMP Users

SNMP users have a specified username, a group to which the user belongs, authentication password, encryption password, and authentication and encryption algorithms to use. The authentication algorithm options are MD5 and SHA. The encryption algorithm options are DES, 3DES, and AES (which is available in 128, 192, and 256 versions). When you create a user, you must associate it with an SNMP group. The user then inherits the security model of the group.

SNMP Hosts

An SNMP host is an IP address to which SNMP notifications and traps are sent. To configure SNMP Version 3 hosts, along with the target IP address, you must configure a username, because traps are only sent to a configured user. SNMP target IP addresses and target parameter names must be unique on the

ASA and ASA Services Module. Each SNMP host can have only one username associated with it. To receive SNMP traps, configure the SNMP NMS, and make sure that you configure the user credentials on the NMS to match the credentials for the ASA and ASASM.

Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software

The SNMP Version 3 implementation in the ASA and ASASM differs from the SNMP Version 3 implementation in the Cisco IOS software in the following ways:

- The local-engine and remote-engine IDs are not configurable. The local engine ID is generated when the ASA or ASASM starts or when a context is created.
- No support exists for view-based access control, which results in unrestricted MIB browsing.
- Support is restricted to the following MIBs: USM, VACM, FRAMEWORK, and TARGET.
- You must create users and groups with the correct security model.
- You must remove users, groups, and hosts in the correct sequence.
- Use of the **snmp-server host** command creates an ASA or ASASM rule to allow incoming SNMP traffic.

Licensing Requirements for SNMP

The following table shows the licensing requirements for this feature:

License Requirement
Base License: Base (DES).
Optional license: Strong (3DES, AES)

Prerequisites for SNMP

SNMP has the following prerequisite:

You must have Cisco Works for Windows or another SNMP MIB-II compliant browser to receive SNMP traps or browse a MIB.

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

- Supported in SNMP Version 3.
- The SNMP client in each ASA or ASASM shares engine data with its peer. Engine data includes the engineID, engineBoots, and engineTime objects of the SNMP-FRAMEWORK-MIB. Engine data is written as a binary file to flash:/snmp/*contextname*.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

- Does not support view-based access control, but the VACM MIB is available for browsing to determine default view settings.
- The ENTITY-MIB is not available in the non-admin context. Use the IF-MIB instead to perform queries in the non-admin context.
- Does not support SNMP Version 3 for the AIP SSM or AIP SSC.
- Does not support SNMP debugging.
- Does not support retireval of ARP information.
- Does not support SNMP SET commands.
- When using NET-SNMP Version 5.4.2.1, only supports the encryption algorithm version of AES128. Does not support the encryption algorithm versions of AES256 or AES192.
- Changes to the existing configuration are rejected if the result places the SNMP feature in an inconsistent state.
- For SNMP Version 3, configuration must occur in the following order: group, user, host.
- Before a group is deleted, you must ensure that all users associated with that group are deleted.
- Before a user is deleted, you must ensure that no hosts are configured that are associated with that username.
- If users have been configured to belong to a particular group with a certain security model, and if the security level of that group is changed, you must do the following in this sequence:
  - Remove the users from that group.
  - Change the group security level.
  - Add users that belong to the new group.
- The creation of custom views to restrict user access to a subset of MIB objects is not supported.
- All requests and traps are available in the default Read/Notify View only.
- The connection-limit-reached trap is generated in the admin context. To generate this trap. you must have at least one snmp-server host configured in the user context in which the connection limit has been reached.
- You cannot query for the chassis temperature for the ASA 5585 SSP-40 (NPE).
- If the NMS cannot successfully request objects or is not correctly handling incoming traps from the ASA, performing a packet capture is the most useful method for determining the problem. Choose **Wizards > Packet Capture Wizard**, and follow the on-screen instructions.

Configuring SNMP

This section describes how to configure SNMP and includes the following topics:

- Enabling SNMP, page 98-6
- Configuring an SNMP Management Station, page 98-6
- Configuring SNMP Traps, page 98-7
- Using SNMP Version 1 or 2c, page 98-8
- Using SNMP Version 3, page 98-8

Enabling SNMP

The SNMP agent that runs on the ASA performs two functions:

- Replies to SNMP requests from NMSs.
- Sends traps (event notifications) to NMSs.

To enable the SNMP agent and identify an NMS that can connect to the SNMP server, see the following pane:

Path	Purpose
Configuration > Device Management > Management Access > SNMP	Ensures that the SNMP server on the ASA or ASASM is enabled. By default, the SNMP server is enabled.

What to Do Next

See the "Configuring an SNMP Management Station" section on page 98-6.

Configuring an SNMP Management Station

To receive requests from the ASA. you must configure an SNMP management station in ASDM.

To configure an SNMP management station, perform the following steps:

Step 1 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 From 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 From 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, from the Username drop-down list, choose the name of a configured user.

Step 9 To specify the method for communicating with this NMS, check either the **Poll** or **Trap** check box.

Step 10 Click **OK**.

The Add SNMP Host Access Entry dialog box closes.

Step 11 Click **Apply**.

The NMS is configured and changes are saved to the running configuration. For more information about SNMP Version 3 NMS tools, see the following URL:

http://www.cisco.com/en/US/docs/security/asa/asa82/snmp/snmpv3_tools.html

What to Do Next

See the "Configuring SNMP Traps" section on page 98-7.

Configuring SNMP Traps

To designate which traps that the SNMP agent generates and how they are collected and sent to NMSs, perform the following steps:

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

Step 2 Click **Configure Traps**.

The SNMP Trap Configuration dialog box appears.

Step 3 The traps are divided into the following categories: standard, IKEv2, entity MIB, IPsec, remote access, resource, NAT, syslog, CPU utilization, CPU utilization and monitoring interval, and SNMP interface threshold. Check the applicable check boxes for the SNMP events to notify through SNMP traps. The default configuration has all SNMP standard traps enabled. If you do not specify a trap type, the default is the syslog trap. The default SNMP traps continue to be enabled with the syslog trap. All other traps are disabled by default. To disable a trap, uncheck the applicable check box. To configure the syslog trap severity level, choose **Configuration > Device Management > Logging > Logging Filters**.

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.

What to Do Next

Choose one of the following:

- See the "Using SNMP Version 1 or 2c" section on page 98-8.
- See the "Using SNMP Version 3" section on page 98-8.

Using SNMP Version 1 or 2c

To configure parameters for SNMP Version 1 or 2c, perform the following steps:

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

Step 2 (Optional) Enter a default community string in the Community String (default) field.

Enter the password used by the SNMP NMSs when sending requests to the ASA. The SNMP community string is a shared secret among the SNMP NMSs and the network nodes being managed. The ASA uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive value up to 32 alphanumeric characters. Spaces are not permitted. The default is public. SNMP Version 2c allows separate community strings to be set for each NMS. If no community string is configured for any NMS, the value set here is used by default.

Step 3 In the Contact field, enter the name of the ASA system administrator. The text is case-sensitive and can be up to 127 alphabetic characters. Spaces are accepted, but multiple spaces are shortened to a single space.

Step 4 In the ASA Location field, enter the location of the ASA 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 ASA port that listens for SNMP requests from NMSs; or keep the default, port number 161.

Step 6 Click **Apply**.

SNMP parameters for Versions 1 and 2c are configured and the changes are saved to the running configuration.

What to Do Next

See the "Monitoring SNMP" section on page 98-9.

Using SNMP Version 3

To configure parameters for SNMP Version 3, perform the following steps:

Step 1 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 After a user has been created, you cannot change the group to which the user belongs.

The Add SNMP User Entry dialog box appears.

Step 3 From the Group Name drop-down list, choose the group to which the SNMP user belongs. 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 Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.

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, choose the 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 alphanumeric 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.

What to Do Next

See the "Monitoring SNMP" section on page 98-9.

Monitoring SNMP

NMSs are the PCs or workstations that you set up to monitor SNMP events and manage devices, such as the ASA.You can monitor the health of a device from an NMS by polling required information from the SNMP agent that has been set up on the device. Predefined events from the SNMP agent to the NMS generate syslog messages. This section includes the following topics:

- SNMP Syslog Messaging, page 98-10
- SNMP Monitoring, page 98-10

SNMP Syslog Messaging

SNMP generates detailed syslog messages that are numbered 212*nnn*. Syslog messages indicate the status of SNMP requests, SNMP traps, SNMP channels, and SNMP responses from the ASA or ASASM to a specified host on a specified interface.

For detailed information about syslog messages, see the syslog messages guide.

Note SNMP polling fails if SNMP syslog messages exceed a high rate (approximately 4000 per second).

SNMP Monitoring

To monitor SNMP, perform the following steps:

Path	Purpose
Tools > Command Line Interface Enter the **show running-config snmp-server** command, then click **Send**.	Shows all SNMP server configuration information.
Tools > Command Line Interface Enter the **show running-config snmp-server group** command, then click **Send**.	Shows SNMP group configuration settings.
Tools > Command Line Interface Enter the **show running-config snmp-server host** command, then click **Send**.	Shows configuration settings used by SNMP to control messages and notifications sent to remote hosts.
Tools > Command Line Interface Enter the **show running-config snmp-server user** command, then click **Send**.	Shows SNMP user-based configuration settings.
Tools > Command Line Interface Type **show snmp-server engineid** command, then click **Send**.	Shows the ID of the SNMP engine configured.
Tools > Command Line Interface Enter the **show snmp-server group** command, then click **Send**.	Shows the names of configured SNMP groups. **Note** If the community string has already been configured, two extra groups appear by default in the output. This behavior is normal.
Tools > Command Line Interface Enter the **show snmp-server statistics** command, then click **Send**.	Shows the configured characteristics of the SNMP server.
Tools > Command Line Interface Enter the **show snmp-server user** command, then click **Send**.	Shows the configured characteristics of users.

Where to Go Next

To configure the syslog server, see Chapter 96, "Configuring Logging."

Additional References

For additional information related to implementing SNMP, see the following sections:

- RFCs for SNMP Version 3, page 98-11
- MIBs, page 98-11
- Application Services and Third-Party Tools, page 98-13

RFCs for SNMP Version 3

RFC	Title
3410	*Introduction and Applicability Statements for Internet Standard Management Framework*
3411	*An Architecture for Describing SNMP Management Frameworks*
3412	*Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
3413	*Simple Network Management Protocol (SNMP) Applications*
3414	*User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMP)*
3826	*The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

MIBs

For a list of supported MIBs and traps for the ASA and ASASM by release, see the following URL:

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Not all OIDs in MIBs are supported. To obtain a list of the supported SNMP MIBs and OIDs for a specific ASA or ASASM, choose **Tools > Command Line Interface**, the following command, then click **Send**:

```
hostname(config)# show snmp-server oidlist
```

Note Although the **oidlist** keyword does not appear in the options list for the **show snmp-server** command help, it is available. However, this command is for Cisco TAC use only. Contact the Cisco TAC before using this command.

The following is sample output from the **show snmp-server oidlist** command:

```
hostname(config)# show snmp-server oidlist
[0]     1.3.6.1.2.1.1.1.        sysDescr
[1]     1.3.6.1.2.1.1.2.        sysObjectID
[2]     1.3.6.1.2.1.1.3.        sysUpTime
[3]     1.3.6.1.2.1.1.4.        sysContact
```

```
[4]     1.3.6.1.2.1.1.5.        sysName
[5]     1.3.6.1.2.1.1.6.        sysLocation
[6]     1.3.6.1.2.1.1.7.        sysServices
[7]     1.3.6.1.2.1.2.1.        ifNumber
[8]     1.3.6.1.2.1.2.2.1.1.    ifIndex
[9]     1.3.6.1.2.1.2.2.1.2.    ifDescr
[10]    1.3.6.1.2.1.2.2.1.3.    ifType
[11]    1.3.6.1.2.1.2.2.1.4.    ifMtu
[12]    1.3.6.1.2.1.2.2.1.5.    ifSpeed
[13]    1.3.6.1.2.1.2.2.1.6.    ifPhysAddress
[14]    1.3.6.1.2.1.2.2.1.7.    ifAdminStatus
[15]    1.3.6.1.2.1.2.2.1.8.    ifOperStatus
[16]    1.3.6.1.2.1.2.2.1.9.    ifLastChange
[17]    1.3.6.1.2.1.2.2.1.10.   ifInOctets
[18]    1.3.6.1.2.1.2.2.1.11.   ifInUcastPkts
[19]    1.3.6.1.2.1.2.2.1.12.   ifInNUcastPkts
[20]    1.3.6.1.2.1.2.2.1.13.   ifInDiscards
[21]    1.3.6.1.2.1.2.2.1.14.   ifInErrors
[22]    1.3.6.1.2.1.2.2.1.16.   ifOutOctets
[23]    1.3.6.1.2.1.2.2.1.17.   ifOutUcastPkts
[24]    1.3.6.1.2.1.2.2.1.18.   ifOutNUcastPkts
[25]    1.3.6.1.2.1.2.2.1.19.   ifOutDiscards
[26]    1.3.6.1.2.1.2.2.1.20.   ifOutErrors
[27]    1.3.6.1.2.1.2.2.1.21.   ifOutQLen
[28]    1.3.6.1.2.1.2.2.1.22.   ifSpecific
[29]    1.3.6.1.2.1.4.1.        ipForwarding
[30]    1.3.6.1.2.1.4.20.1.1.   ipAdEntAddr
[31]    1.3.6.1.2.1.4.20.1.2.   ipAdEntIfIndex
[32]    1.3.6.1.2.1.4.20.1.3.   ipAdEntNetMask
[33]    1.3.6.1.2.1.4.20.1.4.   ipAdEntBcastAddr
[34]    1.3.6.1.2.1.4.20.1.5.   ipAdEntReasmMaxSize
[35]    1.3.6.1.2.1.11.1.       snmpInPkts
[36]    1.3.6.1.2.1.11.2.       snmpOutPkts
[37]    1.3.6.1.2.1.11.3.       snmpInBadVersions
[38]    1.3.6.1.2.1.11.4.       snmpInBadCommunityNames
[39]    1.3.6.1.2.1.11.5.       snmpInBadCommunityUses
[40]    1.3.6.1.2.1.11.6.       snmpInASNParseErrs
[41]    1.3.6.1.2.1.11.8.       snmpInTooBigs
[42]    1.3.6.1.2.1.11.9.       snmpInNoSuchNames
[43]    1.3.6.1.2.1.11.10.      snmpInBadValues
[44]    1.3.6.1.2.1.11.11.      snmpInReadOnlys
[45]    1.3.6.1.2.1.11.12.      snmpInGenErrs
[46]    1.3.6.1.2.1.11.13.      snmpInTotalReqVars
[47]    1.3.6.1.2.1.11.14.      snmpInTotalSetVars
[48]    1.3.6.1.2.1.11.15.      snmpInGetRequests
[49]    1.3.6.1.2.1.11.16.      snmpInGetNexts
[50]    1.3.6.1.2.1.11.17.      snmpInSetRequests
[51]    1.3.6.1.2.1.11.18.      snmpInGetResponses
[52]    1.3.6.1.2.1.11.19.      snmpInTraps
[53]    1.3.6.1.2.1.11.20.      snmpOutTooBigs
[54]    1.3.6.1.2.1.11.21.      snmpOutNoSuchNames
[55]    1.3.6.1.2.1.11.22.      snmpOutBadValues
[56]    1.3.6.1.2.1.11.24.      snmpOutGenErrs
[57]    1.3.6.1.2.1.11.25.      snmpOutGetRequests
[58]    1.3.6.1.2.1.11.26.      snmpOutGetNexts
[59]    1.3.6.1.2.1.11.27.      snmpOutSetRequests
[60]    1.3.6.1.2.1.11.28.      snmpOutGetResponses
[61]    1.3.6.1.2.1.11.29.      snmpOutTraps
[62]    1.3.6.1.2.1.11.30.      snmpEnableAuthenTraps
[63]    1.3.6.1.2.1.11.31.      snmpSilentDrops
[64]    1.3.6.1.2.1.11.32.      snmpProxyDrops
[65]    1.3.6.1.2.1.31.1.1.1.1. ifName
[66]    1.3.6.1.2.1.31.1.1.1.2. ifInMulticastPkts
[67]    1.3.6.1.2.1.31.1.1.1.3. ifInBroadcastPkts
```

```
[68]    1.3.6.1.2.1.31.1.1.1.4. ifOutMulticastPkts
[69]    1.3.6.1.2.1.31.1.1.1.5. ifOutBroadcastPkts
[70]    1.3.6.1.2.1.31.1.1.1.6. ifHCInOctets
--More--
```

Application Services and Third-Party Tools

For information about SNMP support, see the following URL:

http://www.cisco.com/en/US/tech/tk648/tk362/tk605/tsd_technology_support_sub-protocol_home.html

For information about using third-party tools to walk SNMP Version 3 MIBs, see the following URL:

http://www.cisco.com/en/US/docs/security/asa/asa83/snmp/snmpv3_tools.html

Feature History for SNMP

Table 98-2 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 98-2 Feature History for SNMP

Feature Name	Platform Releases	Feature Information
SNMP Versions 1 and 2c	7.0(1)	Provides ASA and ASASM network monitoring and event information by transmitting data between the SNMP server and SNMP agent through the clear text community string. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
SNMP Version 3	8.2(1)	Provides 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 users, groups, and hosts, as well as authentication characteristics by using the USM. In addition, this version allows access control to the agent and MIB objects and includes additional MIB support. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
Password encryption	8.3(1)	Supports password encryption.

Table 98-2 Feature History for SNMP (continued)

Feature Name	Platform Releases	Feature Information
SNMP traps and MIBs	8.4(1)	Supports the following additional keywords: **connection-limit-reached**, **cpu threshold rising**, **entity cpu-temperature**, **entity fan-failure**, **entity power-supply**, **ikev2 stop | start**, **interface-threshold**, **memory-threshold**, **nat packet-discard**, **warmstart**. The entPhysicalTable reports entries for sensors, fans, power supplies, and related components. Supports the following additional MIBs: CISCO-ENTITY-SENSOR-EXT-MIB, CISCO-ENTITY-FRU-CONTROL-MIB, CISCO-PROCESS-MIB, CISCO-ENHANCED-MEMPOOL-MIB, CISCO-L4L7MODULE-RESOURCE-LIMIT-MIB, DISMAN-EVENT-MIB, DISMAN-EXPRESSION-MIB, ENTITY-SENSOR-MIB, NAT-MIB. Supports the following additional traps: ceSensorExtThresholdNotification, clrResourceLimitReached, cpmCPURisingThreshold, mteTriggerFired, natPacketDiscard, warmStart. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
IF-MIB ifAlias OID support	8.2(5)/8.4(2)	The ASA now supports the ifAlias OID. When you browse the IF-MIB, the ifAlias OID will be set to the value that has been set for the interface description.
ASA Services Module (ASASM)	8.5(1)	The ASASM supports all MIBs and traps that are present in 8.4(1), except for the following: Unsupported MIBs in 8.5(1): • CISCO-ENTITY-SENSOR-EXT-MIB (Only objects under the entPhySensorTable group are supported). • ENTITY-SENSOR-MIB (Only objects in the entPhySensorTable group are supported). • DISMAN-EXPRESSION-MIB (Only objects in the expExpressionTable, expObjectTable, and expValueTable groups are supported). Unsupported traps in 8.5(1): • ceSensorExtThresholdNotification (CISCO-ENTITY-SENSOR-EXT-MIB). This trap is only used for power supply failure, fan failure, and high CPU temperature events. • InterfacesBandwidthUtilization.
SNMP traps	8.6(1)	Supports the following additional keywords for the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X: **entity power-supply-presence**, **entity power-supply-failure**, **entity chassis-temperature**, **entity chassis-fan-failure, entity power-supply-temperature**. We modified the following command: **snmp-server enable traps**.

Table 98-2 Feature History for SNMP (continued)

Feature Name	Platform Releases	Feature Information
VPN-related MIBs	9.0(1)	An updated version of the CISCO-IPSEC-FLOW-MONITOR-MIB.my MIB has been implemented to support the next generation encryption feature. The following MIBs have been enabled for the ASASM: • ALTIGA-GLOBAL-REG.my • ALTIGA-LBSSF-STATS-MIB.my • ALTIGA-MIB.my • ALTIGA-SSL-STATS-MIB.my • CISCO-IPSEC-FLOW-MONITOR-MIB.my • CISCO-REMOTE-ACCESS-MONITOR-MIB.my
Cisco TrustSec MIB	9.0(1)	Support for the following MIB was added: CISCO-TRUSTSEC-SXP-MIB.

Configuring Anonymous Reporting and Smart Call Home

The Smart Call Home feature provides personalized, e-mail-based and web-based notification to customers about critical events involving their individual systems, often before customers know that a critical event has occurred.

The Anonymous Reporting feature is a subfeature of the Smart Call Home feature and allows Cisco to anonymously receive minimal error and health information from the device.

Note You might have received a popup dialog that invites you to do the following:

- Enable Anonymous Reporting to help improve the ASA platform.
- Register for Smart Home Notifications to receive personalized, proactive assistance from Cisco.

For information about the dialog, see the "Anonymous Reporting and Smart Call Home Prompt" section on page 99-3.

This chapter describes how to use and configure Anonymous Reporting and Smart Call Home, and it includes the following sections:

- Information About Anonymous Reporting and Smart Call Home, page 99-1
- Licensing Requirements for Anonymous Reporting and Smart Call Home, page 99-4
- Prerequisites for Smart Call Home and Anonymous Reporting, page 99-5
- Guidelines and Limitations, page 99-5
- Configuring Anonymous Reporting and Smart Call Home, page 99-6
- Monitoring Anonymous Reporting and Smart Call Home, page 99-10
- Feature History for Anonymous Reporting and Smart Call Home, page 99-11

Information About Anonymous Reporting and Smart Call Home

This section includes the following topics:

- Information About Anonymous Reporting, page 99-2
- Information About Smart Call Home, page 99-4

Information About Anonymous Reporting

Customers can help to improve the ASA platform by enabling Anonymous Reporting, which allows Cisco to securely receive minimal error and health information from the device. If you enable the feature, your customer identity will remain anonymous, and no identifying information will be sent.

Enabling Anonymous Reporting creates a trust point and installs a certificate. A CA certificate is required for your ASA to validate the server certificate present on the Smart Call Home web server and to form the HTTPS session so that your ASA can send messages securely. Cisco imports a certificate that is predefined in the software. If you decide to enable Anonymous Reporting, a certificate is installed on the ASA with a hardcoded trust point name: _SmartCallHome_ServerCA. When you enable Anonymous Reporting, this trust point is created, the appropriate certificate is installed, and you receive a message about this action. The certificate then appears in your configuration.

If the appropriate certificate already exists in your configuration when you enable Anonymous Reporting, no trust point is created, and no certificate is installed.

Note When you enable Anonymous Reporting, you acknowledge your consent to transfer the specified data to Cisco or to vendors operating on Cisco's behalf (including countries outside of the U.S.). Cisco maintains the privacy of all customers. For information about Cisco's treatment of personal information, see the Cisco Privacy Statement at the following URL: http://www.cisco.com/web/siteassets/legal/privacy.html

What is Sent to Cisco?

Messages are sent to Cisco once a month and whenever the ASA reloads. These messages are categorized by alert groups, which are predefined subsets of Smart Call Home alerts that are supported on the ASA: configuration alerts, inventory alerts, and crash information alerts.

Inventory alerts consist of output from the following commands:

- **show version**—Displays the ASA software version, hardware configuration, license key, and related uptime data for the device.
- **show environment**—Shows system environment information for ASA system components, such as hardware operational status for the chassis, drivers, fans, and power supplies, as well as temperature status, voltage, and CPU usage.
- **show inventory—**Retrieves and displays inventory information about each Cisco product that is installed in the networking device. Each product is identified by unique device information, called the UDI, which is a combination of three separate data elements: the product identifier (PID), the version identifier (VID), and the serial number (SN).
- **show failover state**—Displays the failover state of both units in a failover pair. The information displayed includes the primary or secondary status of the unit, the Active/Standby status of the unit, and the last reported reason for failover.
- **show module**—Shows information about any modules installed on the ASAs, for example, information about an AIP SSC installed on the ASA 5505 or information about an SSP installed on the ASA 5585-X, and information about an IPS SSP installed on an ASA 5585-X.

Configuration alerts consist of output from the following commands:

- **show context**—Shows allocated interfaces and the configuration file URL, the number of contexts configured, or, if you enable Anonymous Reporting in the system execution space, from a list of all contexts.

- **show call-home registered-module status**—Displays the registered module status. If you use system configuration mode, the command displays system module status based on the entire device, not per context.

Upon a system crash, modified information from the following command is sent:

- **show crashinfo** (truncated)—Upon an unexpected software reload, the device sends a modified crash information file with only the traceback section of the file included, so only function calls, register values, and stack dumps are reported to Cisco.

For more information about ASA commands, see the command reference.

DNS Requirement

A DNS server must be configured properly for your ASA to reach the Cisco Smart Call Home server and send messages to Cisco. Because it is possible that your ASA resides in a private network and does not have access to the public network, Cisco verifies your DNS configuration and then configures it for you, if necessary, by doing the following:

1. Performing a DNS lookup for all DNS servers configured.
2. Getting the DNS server from the DHCP server by sending DHCPINFORM messages on the highest security-level interface.
3. Using the Cisco DNS servers for lookup.
4. Randomly using a static IP addresses for tools.cisco.com.

These tasks are performed without changing the current configuration. (For example, the DNS server learned from DHCP will not be added to the configuration.)

If there is no DNS server configured, and your ASA cannot reach the Cisco Smart Call Home Server, Cisco generates a syslog message with the warning severity level for every Smart Call Home message sent to remind you to configure DNS correctly.

For information about syslog messages, see the syslog messages guide.

Anonymous Reporting and Smart Call Home Prompt

When you enter configuration mode you receive a prompt that invites you to enable the Anonymous Reporting and Smart Call Home features if the following criteria are met:

At the prompt, you may choose [Y]es, [N]o, [A]sk later. If you choose [A]sk later, then you are reminded again in seven days or when the ASA reloads. If you continue to choose [A]sk later, the ASA prompts two more times at seven-day intervals before it assumes a [N]o response and does not ask again.

At the ASDM prompt, you can select from the following options:

- Anonymous—Enables Anonymous Reporting.
- Registered (enter an e-mail address)—Enables Smart Call Home and registers your ASA with Cisco TAC.
- Do not enable Smart Call Home—Does not enable Smart Call Home and does not ask again.
- Remind Me Later—Defers the decision. You are reminded again in seven days or whenever the ASA reloads. The ASA prompts two more times at seven-day intervals before it assumes a “Do not enable Smart Call Home response” and does not ask again.

If you did not receive the prompt, you may enable Anonymous Reporting or Smart Call Home by performing the steps in the “Configuring Anonymous Reporting” section on page 99-6 or the “Configuring Smart Call Home” section on page 99-6.

Information About Smart Call Home

When fully configured, Smart Call Home detects issues at your site and reports them back to Cisco or through other user-defined channels (such as e-mail or directly to you), often before you know that these issues exist. Depending upon the seriousness of these problems, Cisco responds to customers regarding their system configuration issues, product end-of-life announcements, security advisory issues, and so on.

In this manner, Smart Call Home offers proactive diagnostics and real-time alerts on the ASA and provides high network availability and increased operational efficiency through proactive and quick issue resolution by doing the following:

- Identifying issues quickly with continuous monitoring, real-time proactive alerts, and detailed diagnostics.
- Making you aware of potential problems through Smart Call Home notifications, in which a service request has been opened, with all diagnostic data attached.
- Resolving critical problems faster with direct, automatic access to experts in Cisco TAC.

Smart Call Home offers increased operational efficiency by providing you with the ability to do the following:

- Use staff resources more efficiently by reducing troubleshooting time.
- Generate service requests to Cisco TAC automatically (if you have a service contract), routed to the appropriate support team, which provides detailed diagnostic information that speeds problem resolution.

The Smart Call Home Portal offers quick, web-based access to required information that provides you with the ability to do the following:

- Review all Smart 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 Smart Call Home-enabled devices.

Licensing Requirements for Anonymous Reporting and Smart Call Home

The following table shows the licensing requirements for Anonymous Reporting and Smart Call Home:

Model	License Requirement
All models	Base License.

Prerequisites for Smart Call Home and Anonymous Reporting

Smart Call Home and Anonymous Reporting have the following prerequisites:

- DNS must be configured. See the "DNS Requirement" section on page 99-3 and the "Configuring the DNS Server" section on page 21-8.

Guidelines and Limitations

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

Context Mode Guidelines

Supported in single mode and multiple context mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines for Anonymous Reporting

- If an Anonymous Reporting message cannot be sent on the first try, the ASA retries two more times before dropping the message.
- Anonymous Reporting can coexist with other Smart Call Home configurations without changing the existing configuration. For example, if Smart Call Home is off before enabling Anonymous Reporting, it remains off, even after enabling Anonymous Reporting.
- If Anonymous Reporting is enabled, you cannot remove the trust point, and when Anonymous Reporting is disabled, the trust point remains. If Anonymous Reporting is disabled, users can remove the trust point, but disabling Anonymous Reporting will not cause the trust point to be removed.

Additional Guidelines for Smart Call Home

- In multiple context mode, the **subscribe-to-alert-group snaphots periodic** command is divided into two commands: one to obtain information from the system configuration and one to obtain information from the user context.
- The Smart Call Home back-end server can accept messages in XML format only.

Configuring Anonymous Reporting and Smart Call Home

While Anonymous Reporting is a subfeature of the Smart Call Home feature and allows Cisco to anonymously receive minimal error and health information from the device, the Smart Call Home feature provides customized support of your system health, enabling Cisco TAC to monitor your devices and open a case when there is an issue, often before you know the issue has occurred.

Generally speaking, you can have both features configured on your system at the same time, yet configuring the Smart Call Home feature provides the same functionality as Anonymous reporting, plus customized services.

This section includes the following topics:

- Configuring Anonymous Reporting, page 99-6
- Configuring Smart Call Home, page 99-6

Configuring Anonymous Reporting

To configure Anonymous Reporting and securely provide minimal error and health information to Cisco, perform the following steps:

Step 1 Choose **Configuration > Device Management > Smart Call Home**.

Step 2 Check the **Enable Anonymous Reporting** check box.

Step 3 Click **Test Connection** to ensure that your system is able to send messages.

ASDM returns a success or error message to notify you of test results.

Step 4 Click **Apply** to save the configuration and enable Anonymous Reporting.

At any time you may also choose to enable the full Smart Call Home feature so that you may receive notification from Cisco about critical events involving your system. You can enable Smart Call Home from the same pane in which you enable Anonymous Reporting. (See the "Monitoring Anonymous Reporting and Smart Call Home" section on page 99-10.)

Configuring Smart Call Home

This section describes how to configure the Smart Call Home feature.

Configuring the Smart Call Home service on your ASA includes the following tasks:

- Enabling the Smart Call Home service.
- Configuring the mail server through which Smart Call Home messages are delivered to subscribers.
- Setting up contact information for the Smart Call Home messages.
- Defining alert processing parameters, such as the maximum rate of events that can be handled.
- Setting up alert subscription profiles. Each alert subscription profile identifies the following:
  - The subscribers to whom the Smart Call Home messages are sent, such as a Smart Call Home server at Cisco or a list of e-mail recipients.
  - Information categories for which you want to receive alerts, such as configuration or inventory information.

Detailed Steps

To configure the Smart Call Home service, system setup, and alert subscription profiles, perform the following steps. To configure Anonymous Reporting, see the "Configuring Anonymous Reporting" section on page 99-6.

Step 1 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.

a. In the Mail Servers pane you can set up mail servers through which Smart Call Home messages are delivered to e-mail subscribers.

b. In the Contact Information pane you can enter the information of the person to contact for the ASA that appears in Smart Call Home messages. This pane includes the following information:

- 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.

c. In the Alert Control pane, administrators can adjust alert control parameters. This pane includes the Alert group status pane, which lists the status (enabled or disabled) of the following alert groups:

- The diagnostics alert group.
- The configuration alert group.
- The environmental alert group.
- The inventory alert group.
- The snapshot alert group.
- The syslog alert group.
- The telemetry alert group.
- 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 a 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 and cannot be edited.

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 which 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, chosen 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 adjust 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 a 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 ASA chooses appropriate values for them.
- For a 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 ASA chooses appropriate values for them.
- For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
- For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.

b. Click the **Basic** or **Detailed** radio button to provide the desired level of information to subscribers.

c. Click **OK** to save the configuration.

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 a 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 ASA chooses appropriate values for them.

- For a 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 ASA chooses appropriate values for them.
- For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
- For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.
- For an interval subscription, specify how often, in minutes, the formation is sent to the subscribers. This requirement is applicable to the snapshot alert group only.

b. Click **OK** to save the configuration.

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** to save the configuration.

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 a 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 ASA chooses appropriate values for them.
- For a 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 ASA chooses appropriate values for them.
- For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
- For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.

b. Click **OK** to save the configuration.

Monitoring Anonymous Reporting and Smart Call Home

To monitor the Anonymous Reporting and Smart Call Home features, navigate to the specified path and enter the specified command:

Path	Purpose
Tools > Command Line Interface Enter the **show call- home detail** command, and click **Send**.	Shows the current Smart Call Home detail configuration.
Tools > Command Line Interface Enter the **show call-home mail-server status** command, and click **Send**.	Shows the current mail server status.
Tools > Command Line Interface Enter the **show smart-call- home profile {***profile name* **| all** command, and click **Send**.	Shows the configuration of Smart Call Home profiles.
Tools > Command Line Interface Enter the **show call-home registered-module all** command, and click **Send**.	Shows the registered module status.
Tools > Command Line Interface Enter the **show smart-call statistics** command, and click **Send**.	Shows call-home detail status.
Tools > Command Line Interface Enter the **show call-home** command, and click **Send**.	Shows the current Smart Call Home configuration.
Tools > Command Line Interface Enter the **show running-config call-home** command, and click **Send**.	Shows the current Smart Call Home running configuration.
Tools > Command Line Interface Enter the **show smart-call-home alert-group** command, and click **Send**.	Shows the current status of Smart Call Home alert groups.
Tools > Command Line Interface Enter the **show running-config all** command, and click **Send**.	Shows details about the Anonymous Reporting user profile.

Feature History for Anonymous Reporting and Smart Call Home

Table 99-1 lists each feature change and the platform release in which it was implemented. ASDM is backward-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 99-1 Feature History for Anonymous Reporting and 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 ASA, and provides higher network availability and increased operational efficiency. We introduced the following screen: Configuration> Device Management> Smart Call Home.
Anonymous Reporting	8.2(5)/8.4(2)	Customers can help to improve the ASA platform by enabling Anonymous Reporting, which allows Cisco to securely receive minimal error and health information from a device. We modified the following screen: Configuration> Device Management> Smart Call Home.

PART 17

System Administration

Managing Software and Configurations

This chapter describes how to manage the ASA software and configurations and includes the following sections:

- Upgrading the Software, page 101-1
- Managing Files, page 101-14
- Configuring the Images and Startup Configuration to Use, page 101-19
- Backing Up and Restoring Configurations or Other Files, page 101-20
- Saving the Running Configuration to a TFTP Server, page 101-27
- Scheduling a System Restart, page 101-27
- Downgrading Your Software, page 101-28
- Configuring Auto Update, page 101-30

Upgrading the Software

Note You can upgrade from any previous release (if available for your model) directly to the latest release. When upgrading to Version 9.0, because of configuration migration, you cannot perform a downgrade; be sure to back up your configuration file in case you want to downgrade.

If you are upgrading to Version 9.0, see the migration section in the release notes for configuration migration information.

If you are upgrading from a pre-8.3 release, see also the *Cisco ASA 5500 Migration Guide to Version 8.3 and Later* for important information about migrating your configuration.

This section describes how to upgrade to the latest version and includes the following topics:

- Viewing Your Current Version, page 101-2
- Downloading the Software from Cisco.com, page 101-2
- Upgrading a Standalone Unit, page 101-2
- Upgrading a Failover Pair or ASA Cluster, page 101-7

Note For CLI procedures, see the ASA documentation.

Viewing Your Current Version

The software version appears on the ASDM home page; view the home page to verify the software version of your ASA.

Downloading the Software from Cisco.com

If you are using the ASDM Upgrade Wizard, you do not have to pre-download the software. If you are manually upgrading, for example for a failover upgrade, download the images to your local computer.

If you have a Cisco.com login, you can obtain the OS and ASDM images from the following website:

http://www.cisco.com/cisco/software/navigator.html?mdfid=279513386

Upgrading a Standalone Unit

Note This section describes how to install the ASDM and operating system (OS) imagesIf the ASA is running Version 8.0 or later, then you can upgrade to the latest version of ASDM (and disconnect and reconnect to start running it) before upgrading the OS. The exception is for ASA versions that are not supported by the latest ASDM version; for example, ASA 8.5. In that case, follow the instructions for pre-8.0 versions (ASDM 5.2 and earlier).

If the ASA is running a version earlier than 8.0, then use the already installed version of ASDM to upgrade both the OS and ASDM to the latest versions, and then reload.

- Upgrading from Your Local Computer (ASDM 6.0 or Later), page 101-2
- Upgrading Using the Cisco.com Wizard (ASDM 6.3 or Later), page 101-4
- Upgrading Using the Cisco.com Wizard (ASDM 6.0 Through ASDM 6.2), page 101-5
- Upgrading from Your Local Computer (ASDM 5.2 or Earlier), page 101-6

Upgrading from Your Local Computer (ASDM 6.0 or Later)

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

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

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 2 In the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.

The Upgrade Software dialog box appears.

Step 3 From the Image to Upload drop-down list, choose **ASDM**.

Step 4 In the Local File Path field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.

Step 5 In the Flash File System Path field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 Repeat Step 2 through Step 6, choosing **ASA** from the Image to Upload drop-down list. You can also use this procedure to upload other file types.

Step 8 Configure the ASA to use the new images.

a. Choose **Configuration > Device Management > System/Image Configuration > Boot Image/Configuration**.

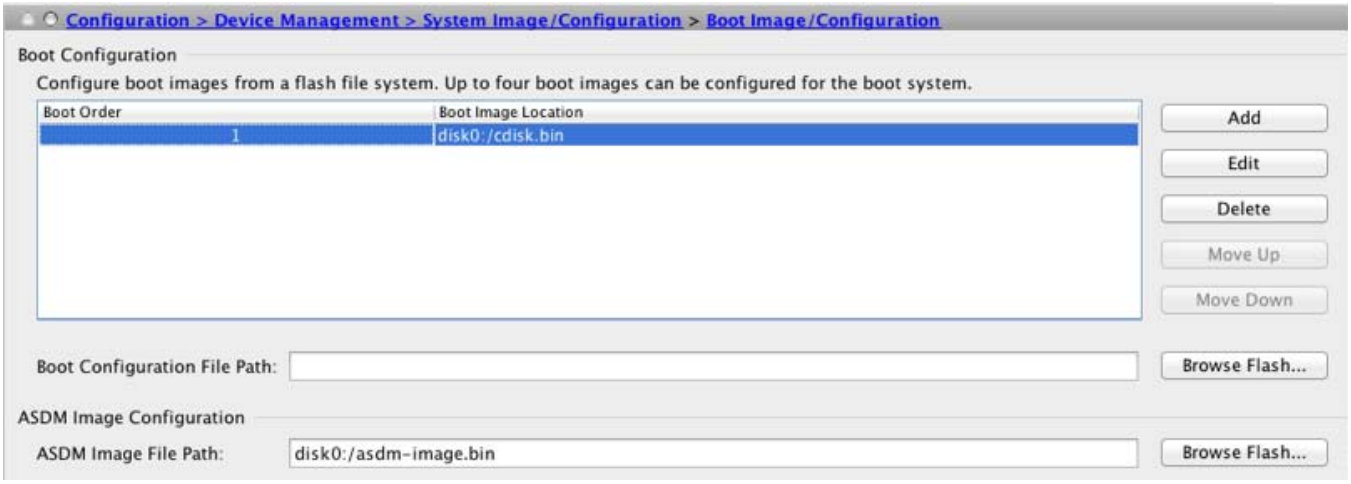

b. In the Boot Configuration table, click **Add** to add the new image (if you have fewer than four images listed); or you can choose an existing image and click **Edit** to change it to the new one.

If you do not specify an image, the ASA searches the internal flash memory for the first valid image to boot; we recommend booting from a specific image.

c. Click **Browse Flash**, choose the OS image, and click **OK**.

d. Click **OK** to return to the Boot Image/Configuration pane.

e. Make sure the new image is the first image in the table by using the **Move Up** button as needed.

f. In the ASDM Image Configuration area, click **Browse Flash**, choose the ASDM image, and click **OK**.

g. Click **Apply**.

Step 9 Choose **File > Save Running Configuration to Flash** to save your configuration changes.

Step 10 Choose **Tools > System Reload** to reload the ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

Step 11 After the ASA reloads, restart ASDM.

Upgrading Using the Cisco.com Wizard (ASDM 6.3 or Later)

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

In this wizard, you can do the following:

- Choose an ASA image file and/or ASDM image file to upgrade.

  Note ASDM downloads the latest image version, which includes the build number. For example, if you are downloading 8.4(2), the dowload might be 8.4(2.8). This behavior is expected, so you may proceed with the planned upgrade.

- Review the upgrade changes that you have made.
- Download the image or images and install them.
- Review the status of the installation.
- If the installation completed successfully, restart the ASA to save the configuration and complete the upgrade.

Detailed Steps

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 2 Choose **Tools > Check for ASA/ASDM Updates**.

In multiple context mode, access this menu from the System.

The Cisco.com Authentication dialog box appears.

Step 3 Enter your assigned Cisco.com username and the Cisco.com password, and then click **Login**.

The Cisco.com Upgrade Wizard appears.

Note If there are no upgrade available, a dialog box appears. Click **OK** to exit the wizard.

Step 4 Click **Next** to display the Select Software screen.

The current ASA version and ASDM version appear.

Step 5 To upgrade the ASA version and ASDM version, perform the following steps:

a. In the ASA area, check the **Upgrade to** check box, and then choose an ASA version to which you want to upgrade from the drop-down list.

b. In the ASDM area, check the **Upgrade to** check box, and then choose an ASDM version to which you want to upgrade from the drop-down list.

Step 6 Click **Next** to display the Review Changes screen.

Step 7 Verify the following items:

- The ASA image file and/or ASDM image file that you have downloaded are the correct ones.
- The ASA image file and/or ASDM image file that you want to upload are the correct ones.
- The correct ASA boot image has been selected.

Step 8 Click **Next** to start the upgrade installation.

You can then view the status of the upgrade installation as it progresses.

The Results screen appears, which provides additional details, such as the upgrade installation status (success or failure).

During the upgrade process from Version 8.2(1) to Version 8.3(1), the following files are automatically saved to flash memory:

- The startup configuration
- The per-context configuration
- The bootup error log, which includes any migration messages

If there is insufficient memory to save the configuration files, an error message appears on the console of the ASA and is saved in the bootup error log file. All previously saved configuration files are also removed.

Step 9 If the upgrade installation succeeded, for the upgrade versions to take effect, check the **Save configuration and reload device now** check box to restart the ASA, and restart ASDM.

Step 10 Click **Finish** to exit the wizard and save the configuration changes that you have made.

Note To upgrade to the next higher version, if any, you must restart the wizard.

Upgrading Using the Cisco.com Wizard (ASDM 6.0 Through ASDM 6.2)

Detailed Steps

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 2 From the Tools menu, choose **Tools > Upgrade Software from Cisco.com**.

In multiple context mode, access this menu from the System.

The Upgrade Software from Cisco.com Wizard appears.

Step 3 Click **Next**.

The Authentication screen appears.

Step 4 Enter your Cisco.com username and password, and click **Next**.

The Image Selection screen appears.

Step 5 Check the **Upgrade the ASA version** check box and the **Upgrade the ASDM version** check box to specify the most current images to which you want to upgrade, and click **Next**.

The Selected Images screen appears.

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. This screen provides additional details, such as whether the upgrade failed or whether you want to save the configuration and reload the ASA.

If you upgraded the ASA version and the upgrade succeeded, an option to save the configuration and reload the ASA appears.

Step 7 Click **Yes**.

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

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

Step 9 After the ASA reloads, restart ASDM.

Upgrading from Your Local Computer (ASDM 5.2 or Earlier)

Detailed Steps

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration. For example, choose **File > Show Running Configuration in New Window** to open the configuration as an HTML page. You can also use one of the File > Save Running Configuration options.

Step 2 Choose **Tools > Upgrade Software**.

Step 3 From the Image to Upload drop-down list, choose **ASDM**.

Step 4 Click **Browse Local Files**, and browse to the ASDM image you downloaded from Cisco.com.

Step 5 Click **Browse Flash** to determine where to install the new ASDM image.

The Browse Flash dialog box appears. Choose the new location, and click **OK**. If you do not have room for both the current image and the new image, you can install over the current image.

Step 6 Click **Upload Image**.

Wait for the image to upload. An information window appears that indicates a successful upload.

Step 7 Repeat Step 2 through Step 6, choosing **ASA** from the Image to Upload drop-down list.

Step 8 Click **Close** to exit the Upgrade Software dialog box.

Step 9 Configure the ASA to use the new images.

a. Choose **Configuration > Properties > Device Administration > Boot Image/Configuration**.

b. In the Boot Configuration table, click **Add** to add the new image (if you have fewer than four images listed); or you can choose an existing image and click **Edit** to change it to the new one.

   If you do not specify an image, the ASA searches the internal flash memory for the first valid image to boot; we recommend booting from a specific image.

c. Click **Browse Flash**, choose the OS image, and click **OK**.

d. Click **OK** to return to the Boot Image/Configuration pane.

e. Make sure the new image is the first image in the table by using the **Move Up** button as needed.

f. In the ASDM Image Configuration area, click **Browse Flash**, choose the ASDM image, and click **OK**.

g. Click **Apply**.

Step 10 Choose **File > Save Running Configuration to Flash** to save your configuration changes.

Step 11 Choose **Tools > System Reload** to reload the ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

Step 12 After the ASA reloads, restart ASDM.

Upgrading a Failover Pair or ASA Cluster

- Software Version Requirements for Zero Downtime Upgrading, page 101-8
- Upgrading an Active/Standby Failover Pair, page 101-8
- Upgrading an ASA Cluster, page 101-12

Software Version Requirements for Zero Downtime Upgrading

The units in a failover configuration or ASA cluster should have the same major (first number) and minor (second number) software version. However, you do not need to maintain version parity on the units during the upgrade process; you can have different versions on the software running on each unit and still maintain failover support. To ensure long-term compatibility and stability, we recommend upgrading all units to the same version as soon as possible.

Table 101-1 shows the supported scenarios for performing zero-downtime upgrades.

Table 101-1 Zero-Downtime Upgrade Support

Type of Upgrade	Support
Maintenance Release	You can upgrade from any maintenance release to any other maintenance release within a minor release. For example, you can upgrade from 7.0(1) to 7.0(4) without first installing the maintenance releases in between.
Minor Release	You can upgrade from a minor release to the next minor release. You cannot skip a minor release. For example, you can upgrade from 7.0(1) to 7.1(1). Upgrading from 7.0(1) directly to 7.2(1) is not supported for zero-downtime upgrades; you must first upgrade to 7.1(1). For models that are not supported on a minor release, you can skip the minor release; for example, for the ASA 5585, you can upgrade from 8.2 to 8.4 (the model is not supported on 8.3).
Major Release	You can upgrade from the last minor release of the previous version to the next major release. For example, you can upgrade from 7.2(1) to 8.0(1), assuming that 7.2(1) is the last minor version in the 7.X release series. **Note** Zero downtime upgrades are possible, even when feature configuration is migrated, for example, from 8.2(X) to 8.3(X).

Upgrading an Active/Standby Failover Pair

To upgrade the Active/Standby failover pair, perform the following steps.

Detailed Steps

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 2 On the active unit, in the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.

The Upgrade Software dialog box appears.

Step 3 From the Image to Upload drop-down list, choose **ASDM**.

Step 4 In the Local File Path field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.

Step 5 In the Flash File System Path field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 Repeat Step 2 through Step 6, choosing **ASA** from the Image to Upload drop-down list.

Step 8 Configure the ASA to use the new images.

a. Choose **Configuration > Device Management > System/Image Configuration > Boot Image/Configuration**.

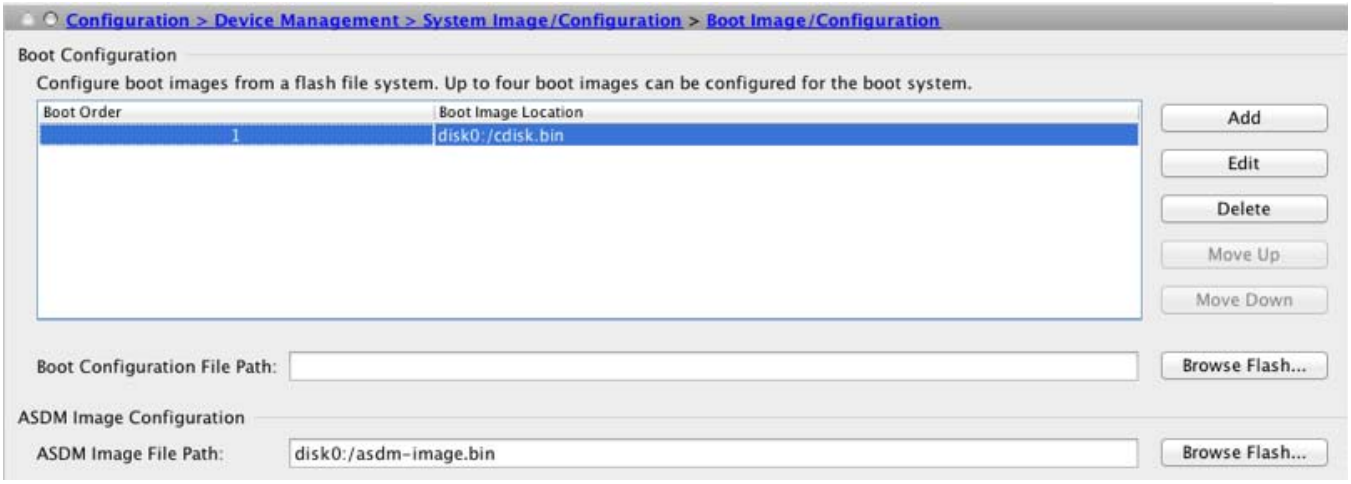

b. In the Boot Configuration table, click **Add** to add the new image (if you have fewer than four images listed); or you can choose an existing image and click **Edit** to change it to the new one.

If you do not specify an image, the ASA searches the internal flash memory for the first valid image to boot; we recommend booting from a specific image.

c. Click **Browse Flash**, choose the OS image, and click **OK**.

d. Click **OK** to return to the Boot Image/Configuration pane.

e. Make sure the new image is the first image in the table by using the **Move Up** button as needed.

f. In the ASDM Image Configuration area, click **Browse Flash**, choose the ASDM image, and click **OK**.

g. Click **Apply**.

Step 9 Choose **File > Save Running Configuration to Flash** to save your configuration changes.

Step 10 Connect ASDM to the *standby* unit, and upload the ASA and ASDM software according to Step 2 through Step 7, using the same file locations you used on the active unit.

Step 11 Choose **Tools > System Reload** to reload the standby ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

Step 12 After the standby ASA reloads, restart ASDM and connect to the standby unit to make sure it is running.

Step 13 Connect ASDM to the *active* unit again.

Step 14 Force the active unit to fail over to the standby unit by choosing **Monitoring > Properties > Failover > Status**, and clicking **Make Standby**.

Step 15 Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

After the ASA comes up, it will now be the standby unit.

Upgrading an Active/Active Failover Pair

To upgrade two units in an Active/Active failover configuration, perform the following steps.

Requirements

Perform these steps in the system execution space.

Detailed Steps

Step 1 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 2 On the primary unit, in the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.

The Upgrade Software dialog box appears.

Step 3 From the Image to Upload drop-down list, choose **ASDM**.

Step 4 In the Local File Path field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.

Step 5 In the Flash File System Path field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 Repeat Step 2 through Step 6, choosing **ASA** from the Image to Upload drop-down list.

Step 8 Configure the ASA to use the new images.

a. Choose **Configuration > Device Management > System/Image Configuration > Boot Image/Configuration**.

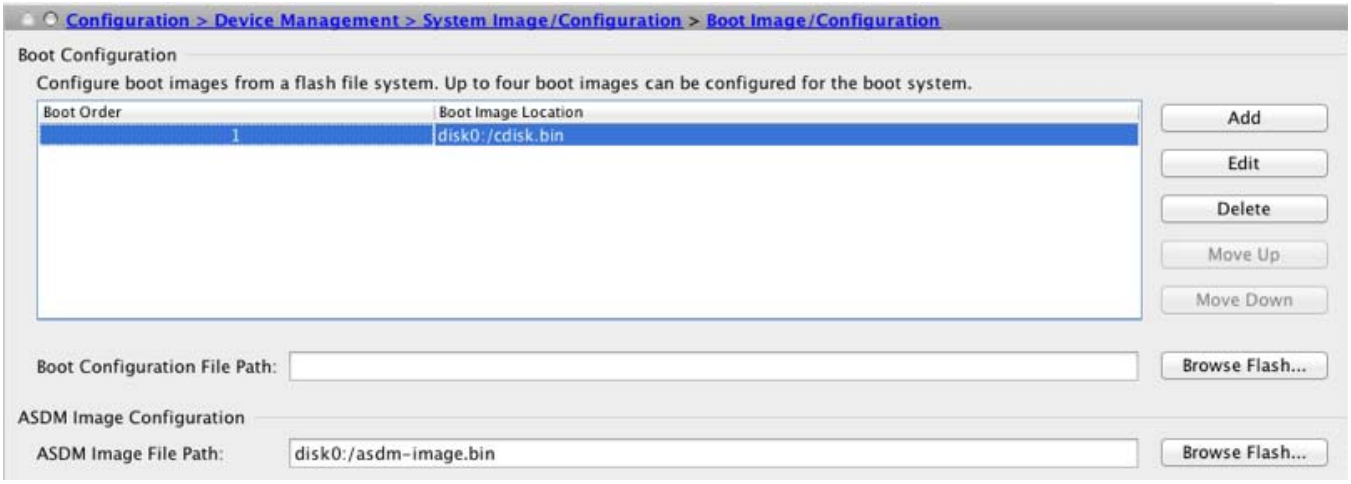

b. In the Boot Configuration table, click **Add** to add the new image (if you have fewer than four images listed); or you can choose an existing image and click **Edit** to change it to the new one.

If you do not specify an image, the ASA searches the internal flash memory for the first valid image to boot; we recommend booting from a specific image.

c. Click **Browse Flash**, choose the OS image, and click **OK**.

d. Click **OK** to return to the Boot Image/Configuration pane.

e. Make sure the new image is the first image in the table by using the **Move Up** button as needed.

f. In the ASDM Image Configuration area, click **Browse Flash**, choose the ASDM image, and click **OK**.

g. Click **Apply**.

Step 9 Choose **File > Save Running Configuration to Flash** to save your configuration changes.

Step 10 Make both failover groups active on the primary unit by choosing **Monitoring > Failover > Failover Group #**, where **#** is the number of the failover group you want to move to the primary unit, and clicking **Make Active**.

Step 11 Connect ASDM to the *secondary* unit, and upload the ASA and ASDM software according to Step 2 through Step 7, using the same file locations you used on the active unit.

Step 12 Choose **Tools > System Reload** to reload the secondary ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

Step 13 Connect ASDM to the *primary* unit, and check when the secondary unit reloads by choosing **Monitoring > Failover > System**.

Step 14 After the secondary unit comes up, force the primary unit to fail over to the secondary unit by choosing **Monitoring > Properties > Failover > System**, and clicking **Make Standby**.

Step 15 Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload. Click the **Save the running configuration at the time of reload** radio button, choose a time to reload (for example, **Now**), and click **Schedule Reload**.

Once the reload is in progress, a Reload Status window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

If the failover groups are configured with Preempt Enabled, they automatically become active on their designated unit after the preempt delay has passed. If the failover groups are not configured with Preempt Enabled, you can return them to active status on their designated units using the **Monitoring > Failover > Failover Group #** pane.

Upgrading an ASA Cluster

To upgrade all units in an ASA cluster, perform the following steps on the master unit. For multiple context mode, perform these steps in the system execution space.

Detailed Steps

Step 1 Launch ASDM on the master unit.

Step 2 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

Step 3 In the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.

The Upgrade Software from Local Computer dialog box appears.

Step 4 Click the **All devices in the cluster** radio button.

Step 5 From the Image to Upload drop-down list, choose the new image file.

Step 6 In the Local File Path field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.

Step 7 In the Flash File System Path field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 8 Click **Upload Image**. The uploading process might take a few minutes; make sure you wait until it is finished.

Step 9 Choose **Tools > System Reload**.

The System Reload dialog box appears.

Step 10 Reload each slave unit one at a time by choosing a slave unit name in the Device drop-down list, and then clicking **Schedule Reload** to reload the unit now.

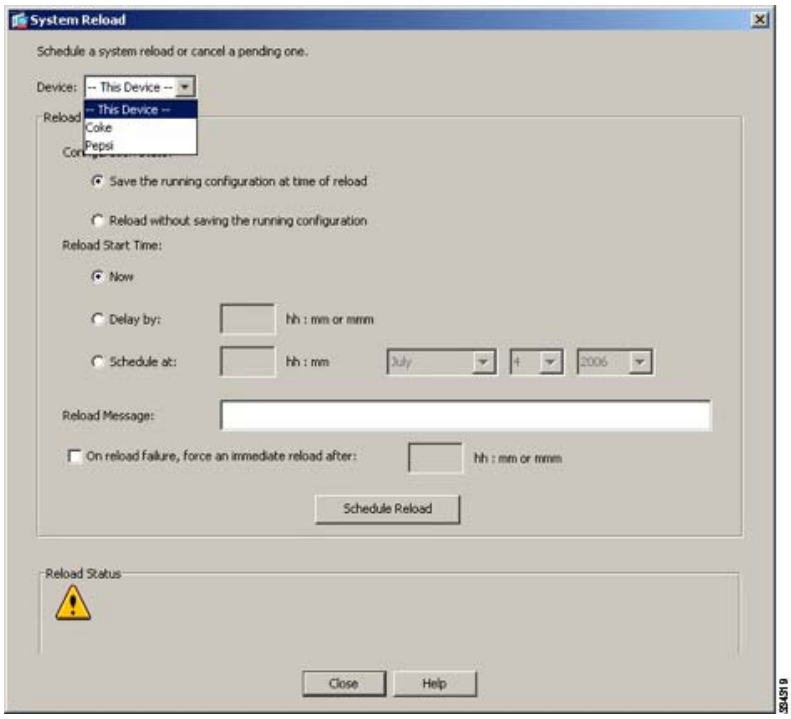

To avoid connection loss, wait for each unit to come back up before reloading the next unit. To view when a unit rejoins the cluster, see the Monitoring > ASA Cluster > Cluster Summary pane.

Step 11 After all slave units have reloaded, reload the master unit from the System Reload dialog box by choosing **--This Device--** from the Device drop-down list.

A new election takes place for a new master unit. When the former master unit rejoins the cluster, it will be a slave.

Step 12 Quit and restart ASDM; you will reconnect to the new master unit.

Managing Files

ASDM provides a set of file management tools to help you perform basic file management tasks. 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.

- Accessing the File Management Tool, page 101-14
- Managing Mount Points, page 101-15
- Transferring Files, page 101-17

Accessing the File Management Tool

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 101-17 for more information.

Step 10 Click **Mount Points** to open the Manage Mount Points dialog box. See the "Managing Mount Points" section on page 101-15 for more information.

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 101-15 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 101-16.

This section includes the following topics:

- Adding or Editing a CIFS/FTP Mount Point, page 101-15
- Accessing a CIFS Mount Point, page 101-16

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 ASA 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 --->
```

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 ASA. You can transfer a remote file to and from the ASA using HTTP, HTTPS, TFTP, FTP, or SMB.

Note For the IPS SSP software module, before you download the IPS software to disk0, make sure at least 50% of the flash memory is free. When you install IPS, IPS reserves 50% of the internal flash memory for its file system.

- Transferring Files Between Local PC and Flash, page 101-18
- Transferring Files Between Remote Server and Flash, page 101-18

Transferring Files Between Local PC and Flash

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.

Transferring Files Between Remote Server and Flash

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 CLI configuration guide.

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.

Configuring the Images and Startup Configuration to Use

By default, the ASA boots the first application image that it finds in internal flash memory. It also boots the first ASDM image it finds in internal flash memory, or if one does not exist in this location, then in external flash memory. If you have more than one image, you should specify the image that you want to boot. For the ASDM image, if you do not specify the image to boot, even if you have only one image installed, then the ASA defines the image in the running configuration. To avoid problems with Auto Update (if configured), and to avoid the image search at each startup, you should specify the ASDM image that you want to boot in the startup configuration.

Choose **Configuration > Device Management > System Image/Configuration > Boot Image/Configuration**.

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. If the device cannot reach the TFTP server to load the image, it tries 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. The Boot Image/Configuration pane includes the following 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. For more information, see the "Adding a Boot Image" section on page 101-20.
- 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 File Path—Displays the location of the configuration file to use at startup.

Adding a 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.

- 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 of the boot image file on the TFTP server, 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.

Backing Up and Restoring Configurations or Other Files

The Backup and Restore features options on the Tools menu let you back up and restore the ASA running configuration, startup configuration, installed add-on images, and SSL VPN Client images and profiles.

The Backup Configurations screen on the ASDM lets you choose the file types to back up, compresses them into a single zip file, then transfer the zip file to the directory that 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.

Note These tools are only available for single context mode.

- Backing Up Configurations, page 101-21
- Backing Up the Local CA Server, page 101-24
- Restoring Configurations, page 101-25
- Saving the Running Configuration to a TFTP Server, page 101-27

Backing Up Configurations

This procedure explains how to back up configurations and images to a .zip file and transfer it to your local computer.

Caution If you have set a master passphrase for the ASA, then you will need that master passphrase to restore the backup configuration that you create with this procedure. If you do not know the master passphrase for the ASA, see the "Configuring the Master Passphrase" section on page 21-4 to learn how to reset it before continuing with the backup.

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

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

The Backup Configurations dialog box appears. Click the down arrow in the **SSL VPN Configuration** area to view the backup options for SSL VPN configurations. By default, all configuration 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.

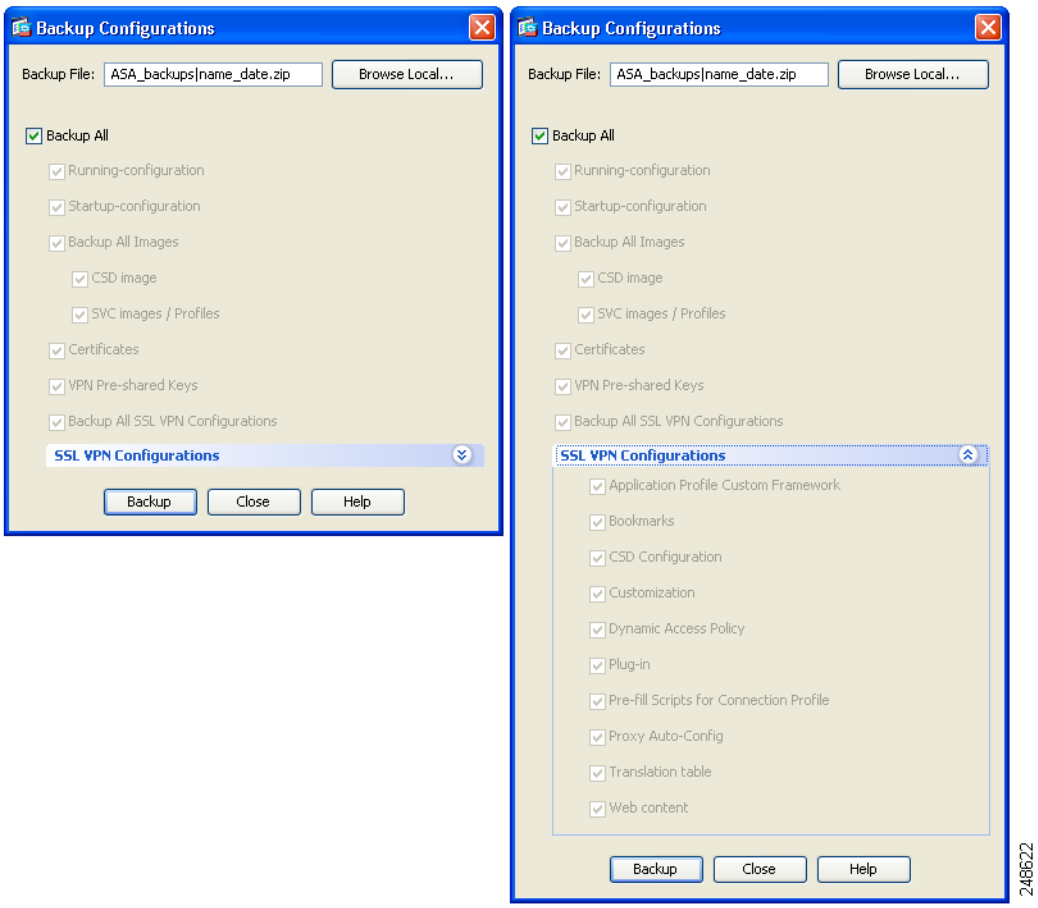

Step 3 Uncheck the **Backup All** check box if you want to select the configurations to back up.

Step 4 Check the check box next to the option that you want to back up.

Step 5 Click **Browse Local to specify a directory and file name for the backup .zip file.**

Step 6 In the Select dialog box, choose the directory in which you want to store the backup file.

Step 7 Click **Select**. The path appears in the Backup File field.

Step 8 Enter the name of the destination backup file after the directory path. The backup file name must be between 3 and 232 characters long.

Step 9 Click **Backup**. The backup proceeds immediately unless you are backing up certificates or the ASA is using a master passphrase.

Step 10 If you have configured and enabled a master passphrase on your ASA, you receive a warning message with a suggestion to change the master passphrase, if you do not know it, before proceeding with the backup. Click **Yes** to proceed with the backup if you know the master passphrase. The backup proceeds immediately unless you are backing up identity certificates.

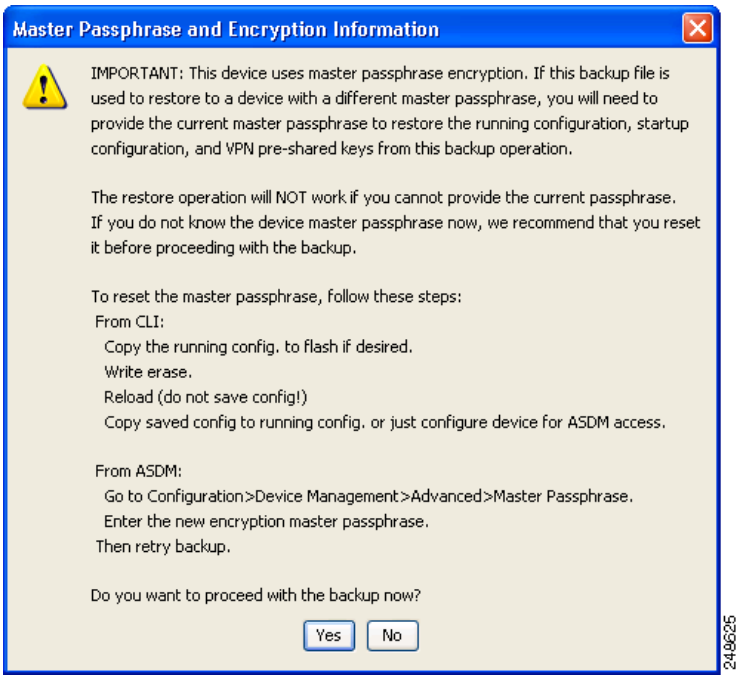

Step 11 If you are backing up an identity certificate, you are asked to enter a separate passphrase to be used for encoding the certificates in PKCS12 format. You can enter a passphrase or skip this step.

Note Identify certificates are backed up by this process; however, certificate authority certificates are not backed up. For instructions on backing up CA certificates, see "Backing Up the Local CA Server" section on page 101-24.

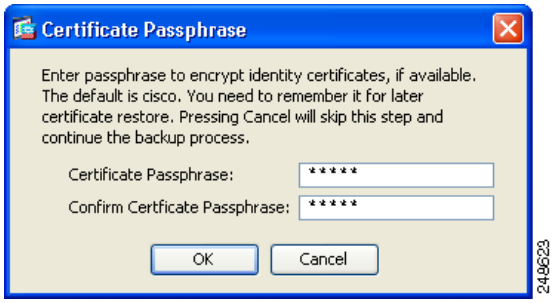

- To encrypt certificates, enter and confirm your certificate passphrase in the Certificate Passphrase dialog box and click **OK**. You will need to remember the password you enter in this dialog box when restoring the certificates.
- Clicking **Cancel** skips the step and does not back up certificates.

After clicking OK or cancel, the backup begins immediately.

Step 12 After the backup is complete, the status window closes and the Backup Statistics dialog box appears to provide success and failure messages.

Note Backup "failure messages" are most likely caused by the lack of an existing configuration for the types indicated.

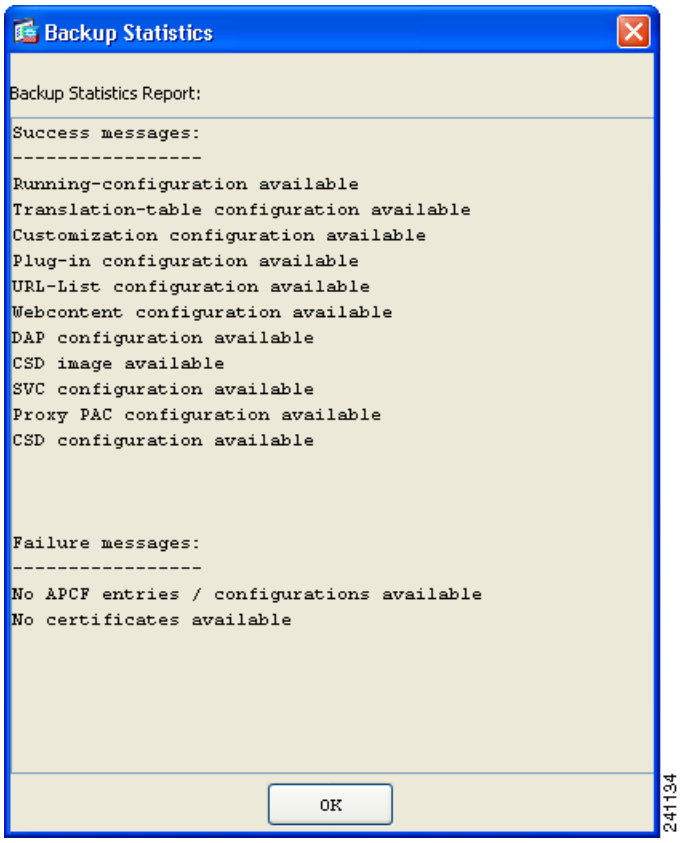

Step 13 Click **OK** to close the Backup Statistics dialog box.

Backing Up the Local CA Server

When you do a ASDM backup, it does not include the local CA server database, so you are not backing up the CA certificates stored on the server. If you want to back up the local CA server, use this manual process with the ASA CLI:

Step 1 Enter the **show run crypto ca server** command.

```
crypto ca server
 keysize server 2048
 subject-name-default OU=aa,O=Cisco,ST=ca,
 issuer-name CN=xxx,OU=yyy,O=Cisco,L=Bxb,St=Mass
 smtp from-address abcd@cisco.com
 publish-crl inside 80
 publish-crl outside 80
```

Step 2 Use the **crypto ca import** command to import the local CA PKCS12 file to create the LOCAL-CA-SERVER trustpoint and to restore the keypair.

```
crypto ca import LOCAL-CA-SERVER pkcs12 <passphrase> (paste the pkcs12
base64 data here)
```

Note Be sure to use the exact name “LOCAL-CA-SERVER” for this step.

Step 3 If the LOCAL-CA-SERVER directory does not exist, you need to create it by entering **mkdir LOCAL-CA-SERVER**.

Step 4 Copy the local CA files into the LOCAL-CA-SERVER directory.

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.ser
disk0:/LOCAL-CA-SERVER/

copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.cdb
disk0:/LOCAL-CA-SERVER/

copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.udb
disk0:/LOCAL-CA-SERVER/

copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.crl
disk0:/LOCAL-CA-SERVER/

copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.p12
disk0:/LOCAL-CA-SERVER/
```

Step 5 Enter the **crypto ca server** command to enable the local CA server

```
crypto ca server
 no shutdown
```

Step 6 Enter the **show crypto ca server** command to check that the local CA server is up and running.

Step 7 Save the configuration.

Restoring Configurations

You can specify configurations and images to restore from a zip file on your local computer.

Before proceeding, note these other restrictions:

- The zip file that you restore must be created by choosing the Tools > Backup Configurations option.
- If you performed the backup with the master passphrase enabled, then you will need that master passphrase in order to restore the running configuration, start-up configuration, and VPN pre-shared keys from the backup you created. If you do not know the master passphrase for the ASA, those items will not be restored during the restore process. See the "Configuring the Master Passphrase" section on page 21-4 for more information on master passphrases.
- If you specified a certificate passphrase during the backup, you will be asked to provide that passphrase in order to restore the certificates. The default passphrase is `cisco`.
- 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 ASA type. You must start with a basic configuration that allows ASDM access.

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

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

Step 2 In the Restore Configurations dialog box, click **Browse Local Directory**, choose the zip file on your local computer that contains the configuration to restore, then click **Select**. The path and the zip filename appear in the Local File field.

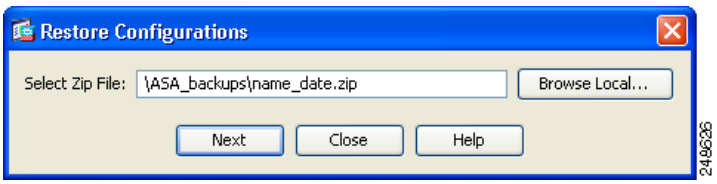

Step 3 Click **Next**. The second Restore Configuration dialog box appears. Check the check boxes next to the configurations that you want to restore. All available SSL VPN configurations are selected by default.

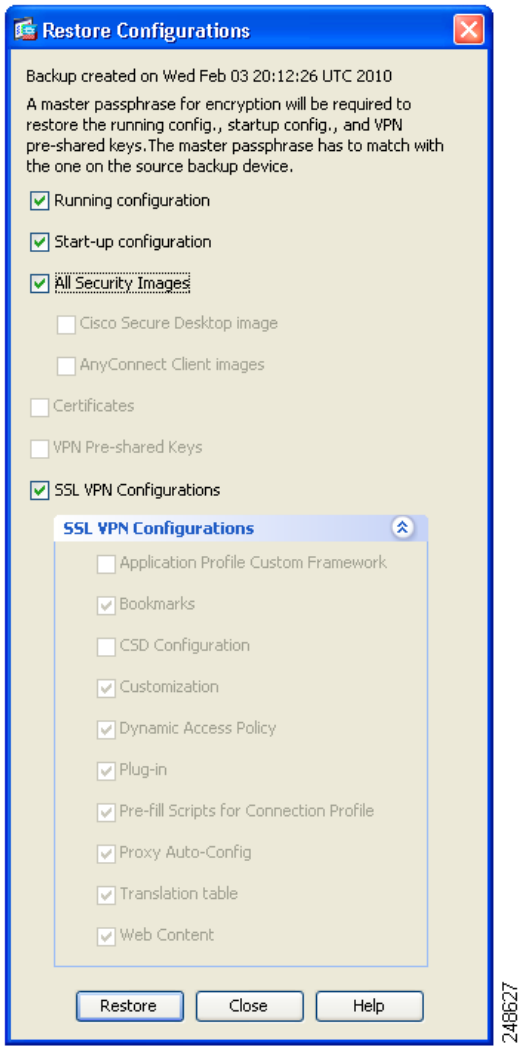

Step 4 Click **Restore**.

Step 5 If you specified a certificate passphrase with which to encrypt the certificates when you created the backup file, ASDM prompts you to enter the passphrase.

Step 6 If you chose to restore the running configuration, you are asked if you want to merge the running configuration, replace the running configuration, or skip this part of the restoration process.

- Merging configurations combines the current running configuration and the backed-up running configuration.

- Replacing the running configuration uses the backed-up running configuration only.
- Skipping the step does not restore the backed-up running configuration.

ASDM displays a status dialog box until the restore operation is finished.

Step 7 If you replaced or merged the running configuration, close ASDM and restart it. If you did not restore the running configuration or the running configuration, refresh the ASDM session for the changes to take effect.

Saving 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.

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.

Downgrading Your Software

When you upgrade to Version 8.3, your configuration is migrated. The old configuration is automatically stored in flash memory. For example, when you upgrade from Version 8.2(1) to 8.3(1), the old 8.2(1) configuration is stored in flash memory in a file called 8_2_1_0_startup_cfg.sav.

Note You must manually restore the old configuration before downgrading.

This section describes how to downgrade and includes the following topics:

- Information About Activation Key Compatibility, page 101-28
- Performing the Downgrade, page 101-29

Information About Activation Key Compatibility

Your activation key remains compatible if you upgrade to the latest version from any previous version. However, you might have issues if you want to maintain downgrade capability:

- Downgrading to Version 8.1 or earlier versions—After you upgrade, if you activate additional feature licenses that were introduced *before 8.2*, the activation key continues to be compatible with earlier versions if you downgrade. However if you activate feature licenses that were introduced in Version 8.2 or later versions, the activation key is not backwards compatible. If you have an incompatible license key, see the following guidelines:
  - If you previously entered an activation key in an earlier version, the ASA uses that key (without any of the new licenses you activated in Version 8.2 or later versions).
  - If you have a new system and do not have an earlier activation key, you need to request a new activation key compatible with the earlier version.
- Downgrading to Version 8.2 or earlier versions—Version 8.3 introduced more robust time-based key usage as well as failover license changes:
  - If you have more than one time-based activation key active, when you downgrade, only the most recently activated time-based key can be active. Any other keys are made inactive.
  - If you have mismatched licenses on a failover pair, downgrading will disable failover. Even if the keys are matching, the license used will no longer be a combined license.

Performing the Downgrade

See the "The Backup and Restore features options on the Tools menu let you back up and restore the ASA running configuration, startup configuration, installed add-on images, and SSL VPN Client images and profiles." section on page 101-20 for more information about configuration migration.

To downgrade from Version 8.3, perform the following steps:

Detailed Steps

Step 1 Choose **Tools > Downgrade Software**.

The Downgrade Software dialog box appears.

Figure 101-1 Downgrade Software

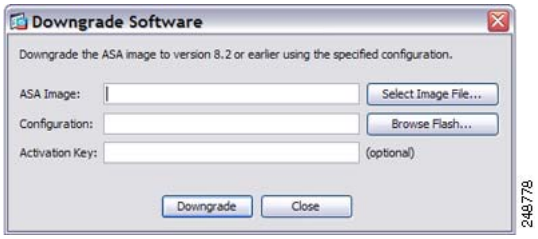

Step 2 For the ASA Image, click **Select Image File**.

The Browse File Locations dialog box appears.

Step 3 Click one of the following radio buttons:

- **Remote Server**—Choose **ftp**, **smb**, or **http** from the drop-down list, and type the path to the old image file.
- **Flash File System**—Click **Browse Flash** to choose the old image file on the local flash file system.

Step 4 For the Configuration, click **Browse Flash** to choose the pre-migration configuration file. (By default this was saved on disk0).

Step 5 (Optional) In the Activation Key field, enter the old activation key if you need to revert to a pre-8.3 activation key.

See the "Information About Activation Key Compatibility" section on page 101-28 for more information.

Step 6 Click **Downgrade**.

This tool is a shortcut for completing the following functions:

1. Clearing the boot image configuration (**clear configure boot**).
2. Setting the boot image to be the old image (**boot system**).
3. (Optional) Entering a new activation key (**activation-key**).
4. Saving the running configuration to startup (**write memory**). This sets the BOOT environment variable to the old image, so when you reload, the old image is loaded.
5. Copying the old configuration to the startup configuration (**copy** *old_config_url* **startup-config**).
6. Reloading (**reload**).

Configuring Auto Update

This section includes the following topics:

- Information About Auto Update, page 101-30
- Configuring Communication with an Auto Update Server, page 101-31

Information About Auto Update

Auto Update is a protocol specification that allows an Auto Update Server to download configurations and software images to many ASAs and can provide basic monitoring of the ASAs from a central location.

The ASA can be configured as either a client or a server. As an Auto Update client, it periodically polls the Auto Update Server for updates to software images and configuration files. As an Auto Update Server, it issues updates for ASAs configured as Auto Update clients.

Auto Update is useful in solving many issues facing administrators for ASA management, such as:

- Overcoming dynamic addressing and NAT challenges.
- Committing configuration changes in one action.
- Providing a reliable method for updating software.
- Leveraging well-understood methods for high availability (failover).
- Providing flexibility with an open interface.
- Simplifying security solutions for Service Provider environments.

The Auto Update specification provides the infrastructure necessary for remote management applications to download ASA 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 ASA, or to pull configuration information by having the ASA periodically poll the Auto Update server. The Auto Update server can also send a command to the ASA to send an immediate polling request at any time. Communication between the Auto Update server and the ASA requires a communications path and local CLI configuration on each ASA.

Guidelines and Limitations

- If the ASA 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 obtain 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 ASA uses SSL, which requires the ASA to have a DES or 3DES license.
- Auto Update is supported in single context mode only.

Configuring Communication with an Auto Update Server

Detailed Steps

To configure the Auto Update feature, choose **Configuration > Device Management > System Image/Configuration > Auto Update**. 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 ASA polls the server listed at the top of the table first. To change the order of the servers in the table, click **Move Up** or **Move Down**. The Auto Update Servers table includes 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 ASA checks the certificate returned by the Auto Update server with the CA root certificates. The Auto Update server and the ASA must 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 ASA waits for the Auto Update server to time out. The Timeout area includes the following fields:

- Enable Timeout Period—Check to enable the ASA to time out if no response is received from the Auto Update server.
- Timeout Period (Minutes)—Enter the number of minutes the ASA will wait to time out if no response is received from the Auto Update server.

The Polling area lets you configure how often the ASA will poll for information from the Auto Update server. The Polling area includes the following fields:

- Polling Period (minutes)—The number of minutes the ASA 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 ASA 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 ASA will attempt to retry to poll the Auto Update server for new information.

Adding or Editing an Auto Update Server

The Add/Edit Auto Update Server dialog box includes the following fields:

- URL—The protocol that the Auto Update server uses to communicate with the ASA, 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—Click to enable the ASA to verify the certificate returned by the Auto Update server with the CA root certificates. The Auto Update server and the ASA must use the same CA.

The User area includes 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.
- Use Device ID to uniquely identify the ASA—Enables authentication using a device ID. The device ID is used to uniquely identify the ASA to the Auto Update server.
- Device ID—Type of device ID to use.
  - Hostname—The name of the host.
  - Serial Number—The device serial number.
  - IP Address on interface—The IP address of the selected interface, used to uniquely identify the ASA to the Auto Update server.
  - MAC Address on interface—The MAC address of the selected interface, used to uniquely identify the ASA to the Auto Update server.
  - User-defined value—A unique user ID.

Setting the Polling Schedule

The Set Polling Schedule dialog box lets you configure specific days and the time-of-day for the ASA to poll the Auto Update server.

The Set Polling Schedule dialog box includes the following fields:

Days of the Week—Check the days of the week that you want the ASA to poll the Auto Update server.

The Daily Update pane group lets you configure the time of day when you want the ASA to poll the Auto Update server, and includes the following fields:

- Start Time—Enter the hour and minute to begin the Auto Update poll.
- Enable randomization—Check to enable the ASA to randomly choose a time to poll the Auto Update server.

Troubleshooting

This chapter describes how to troubleshoot the ASA and includes the following sections:

- Configuring and Running Captures with the Packet Capture Wizard, page 102-1

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. Those tasks include the following:

- Selecting an ingress interface.
- Selecting an egress interface.
- Setting the buffer parameters.
- Running the captures.
- Saving the captures to your PC (optional).

Step 2 Click **Next**.

In a clustering environment, the Cluster Option screen appears. Go to Step 3.

Note For more information about clustering, see Chapter 10, “Configuring a Cluster of ASAs.”

In a non-clustering environment, the Ingress Traffic Selector screen appears. Go to Step 4.

Step 3 In the Cluster Option screen, choose one of the following options for running a capture: **This device only** or **The whole cluster**, then click **Next**.

Step 4 Choose the ingress interface from the drop-down list.

Step 5 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 6 Click **Next** to display the Ingress Traffic Selector screen. For more information, see the "Ingress Traffic Selector" section on page 102-3.

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

Step 8 Enter the destination host IP address and choose the network IP address from the drop-down list.

Step 9 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 10 Click **Next** to display the Egress Traffic Selector screen. For more information, see the "Egress Traffic Selector" section on page 102-4.

Step 11 Choose the egress interface from the drop-down list.

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

Step 13 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 and are based on the choices that you made in the Ingress Traffic Selector screen.

Step 14 Click **Next** to display the Buffers & Captures screen. For more information, see the "Buffers" section on page 102-4.

Step 15 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 16 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 17 Click **Next** to display the Summary screen, which shows the cluster options for all units in the cluster (if you are using clustering), traffic selectors, and buffer parameters that you have entered. For more information, see the "Summary" section on page 102-4.

Step 18 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 102-4. If you are using clustering, go to Step 20.

Step 19 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 20 In a clustering environment, on the Run Captures screen, perform one or more of the following steps:

- Click **Get Cluster Capture Summary** to view a summary of packet capture information for all units in the cluster, followed by packet capture information for each unit.
- Click **Get Capture Buffer** to determine how much buffer space you have remaining in each unit of the cluster. The Capture Buffer from Device dialog box appears.
- Click **Clear Buffer on Device** to remove the current content for one or all of the units in a cluster and allow room in the buffer to capture more packets.

Step 21 Click **Save captures** to display the Save Capture dialog box. You have the option of saving either the ingress capture, the egress capture, or both. For more information, see the "Save Captures" section on page 102-5.

Step 22 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 23 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the ingress capture.

Step 24 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 25 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the egress capture.

Step 26 Click **Close,** then click **Finish** to exit the wizard.

Ingress Traffic Selector

To configure the ingress interface, source and destination hosts or 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 are ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.

a. Enter the ICMP type for ICMP 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.

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.

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.

Summary

The Summary screen shows the cluster options (if you are using clustering), traffic selectors, and the buffer parameters for the packet capture selected in the previous wizard screens.

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**.

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**.

PART 18

Reference

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 A-1
- IPv6 Addresses, page A-5
- Protocols and Applications, page A-11
- TCP and UDP Ports, page A-11
- Local Ports and Protocols, page A-14
- ICMP Types, page A-15

IPv4 Addresses and Subnet Masks

This section describes how to use IPv4 addresses in the ASA. 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 A-1
- Private Networks, page A-2
- Subnet Masks, page A-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, then 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 A-3
- Determining the Address to Use with the Subnet Mask, page A-3

Determining the Subnet Mask

To determine the subnet mask based on how many hosts you want, see Table A-1.

Table A-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 A-3
- Class B-Size Network Address, page A-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, Table A-2 shows the 8-host subnets (/29) of 192.168.0.x.

Table A-2 Class C-Size Network Address

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

Table A-2 Class C-Size Network Address (continued)

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, Table A-3 shows the 16 subnets of the network 10.1.

Table A-3 Subnets of Network

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 A-5
- IPv6 Address Types, page A-6
- IPv6 Address Prefixes, page A-10

Note This section describes the IPv6 address format, the types, and prefixes. For information about configuring the ASA to use IPv6, see the "Configuring IPv6 Addressing" section on page 17-15

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.

You do not need 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 A-4 shows several examples of address compression for different types of IPv6 address.

Table A-4 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 A-6
- Multicast Address, page A-8
- Anycast Address, page A-9
- Required Addresses, page A-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 A-7
- Site-Local Address, page A-7
- Link-Local Address, page A-7
- IPv4-Compatible IPv6 Addresses, page A-7
- Unspecified Address, page A-8
- Loopback Address, page A-8
- Interface Identifiers, page A-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 the "IPv6 Address Prefixes" section on page A-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 the "Interface Identifiers" section on page A-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 the "IPv4-Compatible IPv6 Addresses" section on page A-7).

Site-Local Address

Site-local addresses are used for addressing within a site. They can be used 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 A-1 shows the format of the IPv6 multicast address.

Figure A-1 IPv6 Multicast Address Format

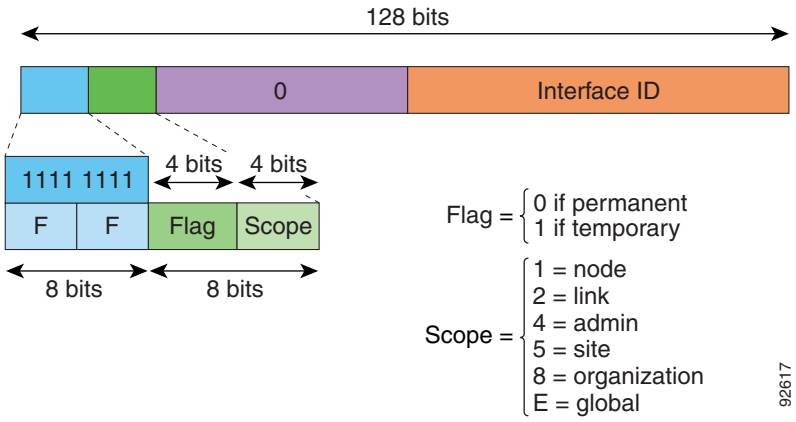

IPv6 nodes (hosts and routers) are required to join the following multicast groups:

- 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 ASA.

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 A-5 shows the prefixes for each IPv6 address type.

Table A-5 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 A-6 lists the protocol literal values and port numbers; either can be entered in ASA commands.

Table A-6 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 A-7 lists the literal values and port numbers; either can be entered in ASA commands. See the following caveats:

- The ASA 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 ASA listens for RADIUS on ports 1645 and 1646. If your RADIUS server uses the standard ports 1812 and 1813, you can configure the ASA 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 ASA 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 A-7 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
isakmp	UDP	500	Internet Security Association and Key Management Protocol
kerberos	TCP, UDP	750	Kerberos

Table A-7 Port Literal Values (continued)

Literal	TCP or UDP?	Value	Description
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
tacacs	TCP, UDP	49	Terminal Access Controller Access Control System Plus
talk	TCP, UDP	517	Talk
telnet	TCP	23	RFC 854 Telnet

Table A-7 Port Literal Values (continued)

Literal	TCP or UDP?	Value	Description
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 A-8 lists the protocols, TCP ports, and UDP ports that the ASA may open to process traffic destined to the ASA. Unless you enable the features and services listed in Table A-8, the ASA does *not* open any local protocols or any TCP or UDP ports. You must configure a feature or service for the ASA 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 A-8 Protocols and Ports Opened by Features and Services

Feature or Service	Protocol	Port Number	Comments
DHCP	UDP	67,68	—
Failover Control	105	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.
NTP	UDP	123	—
OSPF	89	N/A	Protocol only open on destination IP address 224.0.0.5 and 224.0.0.6

Table A-8 Protocols and Ports Opened by Features and Services (continued)

Feature or Service	Protocol	Port Number	Comments
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	8 (non-secure) 9 (secure)	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 A-9 lists the ICMP type numbers and names that you can enter in ASA commands.

Table A-9 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
18	mask-reply

Table A-9 ICMP Types (continued)

ICMP Number	ICMP Name
31	conversion-error
32	mobile-redirect

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 ASA. Before you configure the ASA to use an external server, you must configure the server with the correct ASA 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, page B-1
- Configuring an External LDAP Server, page B-2
- Configuring an External RADIUS Server, page B-25
- Configuring an External TACACS+ Server, page B-37

Understanding Policy Enforcement of Permissions and Attributes

The ASA supports several methods of applying user authorization attributes (also called user entitlements or permissions) to VPN connections. You can configure the ASA to obtain user attributes from a Dynamic Access Policy (DAP) on the ASA, from an external authentication and/or authorization AAA server (RADIUS or LDAP), from a group policy on the ASA, or from all three.

If the ASA 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 ASA applies attributes in the following order (see Figure B-1).

1. DAP attributes on the ASA—Introduced in Version 8.0(2), these attributes take precedence over all others. If you set a bookmark or URL list in DAP, it overrides a bookmark or URL list set in the group policy.
2. User attributes on the AAA server—The server returns these attributes 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 ASA (User Accounts in ASDM).
3. Group policy configured on the ASA—If a RADIUS server returns the value of the RADIUS CLASS attribute IETF-Class-25 (OU=*group-policy*) for the user, the ASA 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 that you configure on the ASA maps the LDAP attribute to the Cisco attribute IETF-Radius-Class.

4. Group policy assigned by the Connection Profile (called tunnel-group in the 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 ASA 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 ASA (DfltGrpPolicy)—System default attributes provide any values that are missing from the DAP, user attributes, group policy, or connection profile.

Figure B-1 Policy Enforcement Flow

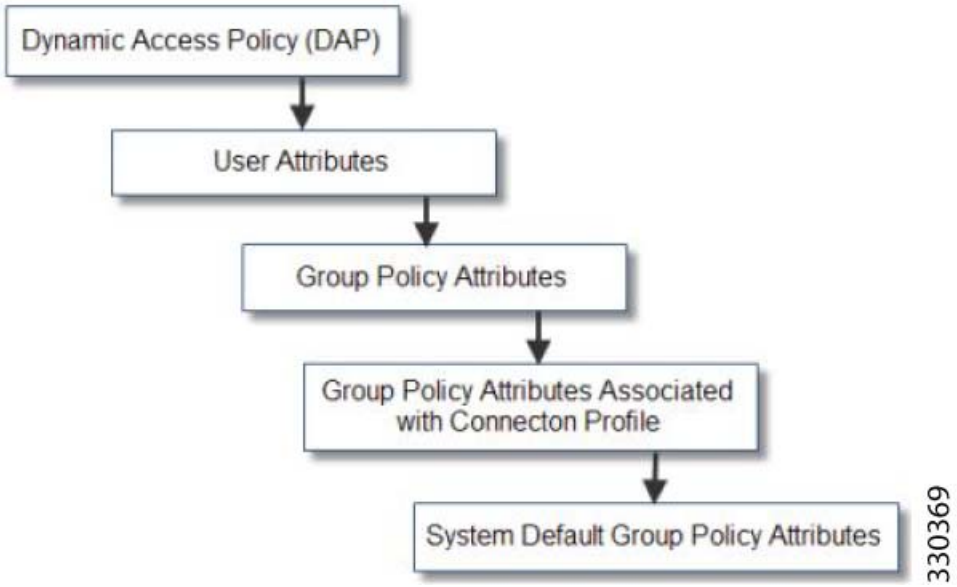

Configuring an External LDAP Server

The VPN 3000 concentrator and the ASA/PIX 7.0 software required a Cisco LDAP schema for authorization operations. Beginning with Version 7.1.x, the ASA 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 the "Active Directory/LDAP VPN Remote Access Authorization Examples" section on page B-15.

This section describes the structure, schema, and attributes of an LDAP server and includes the following topics:

- Organizing the ASA for LDAP Operations, page B-3
- Defining the ASA LDAP Configuration, page B-5
- Active Directory/LDAP VPN Remote Access Authorization Examples, page B-15

The specific steps of these processes vary, depending on which type of LDAP server that you are using.

Note For more information about the LDAP protocol, see RFCs 1777, 2251, and 2849.

Organizing the ASA for LDAP Operations

This section describes how to search within the LDAP hierarchy and perform authenticated binding to the LDAP server on the ASA and includes the following topics:

- Searching the LDAP Hierarchy, page B-3
- Binding the ASA to the LDAP Server, page B-4

Your LDAP configuration should reflect the logical hierarchy of your organization. For example, suppose an employee at your company, Example Corporation, is named Employee1. Employee1 works in the Engineering group. Your LDAP hierarchy could have one or many levels. You might decide to set up a single-level hierarchy in which Employee1 is considered a member of Example Corporation. Or you could set up a multi-level hierarchy in which Employee1 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 B-2 for an example of a multi-level hierarchy.

A multi-level hierarchy has more detail, but searches return results more quickly in a single-level hierarchy.

Figure B-2 A Multi-Level LDAP Hierarchy

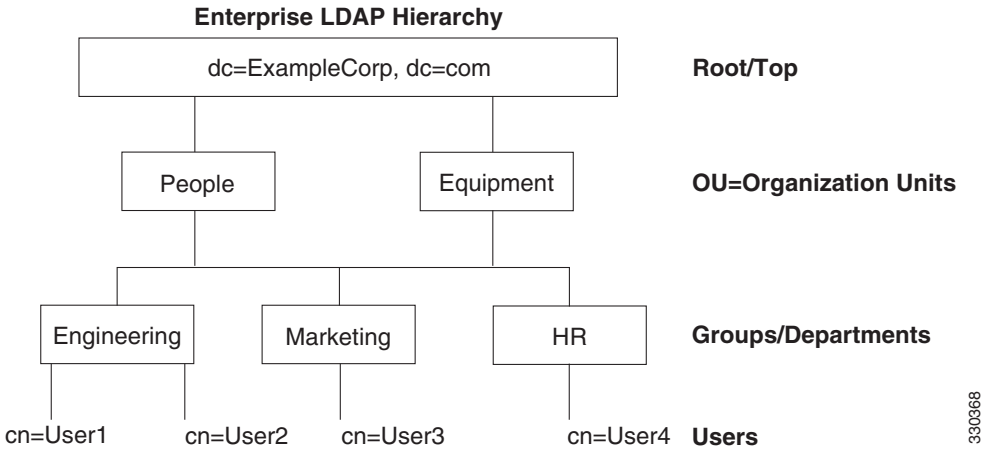

Searching the LDAP Hierarchy

The ASA lets you tailor the search within the LDAP hierarchy. You configure the following three fields on the ASA to define where in the LDAP hierarchy that 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 that includes the user permissions.

- LDAP Base DN defines where in the LDAP hierarchy that the server should begin searching for user information when it receives an authorization request from the ASA.

- 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 it, 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 B-2 shows a sample LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. Table B-1 shows two sample search configurations.

In the first example configuration, when Employee1 establishes the IPsec tunnel with LDAP authorization required, the ASA sends a search request to the LDAP server, indicating it should search for Employee1 in the Engineering group. This search is quick.

In the second example configuration, the ASA sends a search request indicating that the server should search for Employee1 within Example Corporation. This search takes longer.

Table B-1 Example Search Configurations

No.	LDAP Base DN	Search Scope	Naming Attribute	Result
1	group= Engineering,ou=People,dc=ExampleCorporation, dc=com	One Level	cn=Employee1	Quicker search
2	dc=ExampleCorporation,dc=com	Subtree	cn=Employee1	Longer search

Binding the ASA to the LDAP Server

Some LDAP servers (including the Microsoft Active Directory server) require the ASA to establish a handshake via authenticated binding before they accept requests for any other LDAP operations. The ASA uses the Login Distinguished Name (DN) and Login Password to establish a trust relationship (bind) with an LDAP server before a user can search. The Login DN represents a user record in the LDAP server that the administrator uses for binding.

When binding, the ASA 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 ASA can bind with a Login DN with fewer privileges. For example, the Login DN can be a user whose AD “Member Of” designation is part of Domain Users. For VPN password management write operations, the Login DN needs elevated privileges and must be part of the Account Operators AD group. Microsoft Active Directory group search (also called “MemberOf retrieval”) was added in ASA Version 8.0.4.

An example of a Login DN includes the following entries:

cn=Binduser1,ou=Admins,ou=Users,dc=company_A,dc=com

See your LDAP Administrator guide for specific Login DN requirements for read and write operations.

The ASA supports the following features:

- Simple LDAP authentication with an unencrypted password using the default port 389 . You can also use other ports instead of the default port.
- Secure LDAP (LDAP-S) using the default port 636. You can also use other ports instead of the default port.
- Simple Authentication and Security Layer (SASL) MD5
- SASL Kerberos

The ASA does not support anonymous authentication.

Note As an LDAP client, the ASA does not support the transmission of anonymous binds or requests.

Defining the ASA LDAP Configuration

This section describes how to define the LDAP AV-pair attribute syntax and includes the following topics:

- Supported Cisco Attributes for LDAP Authorization, page B-5
- Cisco AV Pair Attribute Syntax, page B-12
- Cisco AV Pairs ACL Examples, page B-13

Note The ASA 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 enforces permissions or attributes if they are configured.

For software Version 7.0, LDAP attributes include the cVPN3000 prefix. For software Versions 7.1 and later, this prefix was removed.

Supported Cisco Attributes for LDAP Authorization

This section provides a complete list of attributes (see Table B-2) for the ASA 5500, VPN 3000 concentrator, and PIX 500 series ASAs. The table includes attribute support information for the VPN 3000 concentrator and PIX 500 series ASAs to assist you in configuring networks with a combination of these devices.

Table B-2 ASA 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 B-2 ASA 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 the "Cisco AV Pair Attribute Syntax" section on page B-12."
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, and 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 B-2 ASA 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 ASA 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 versions 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 ASA. The setting applies to VPN remote access IPsec and SSL VPN clients.
IETF-Radius-Framed-IP-Address	Y	Y	Y	String	Single	An IP address. The setting applies to VPN remote access IPsec and SSL VPN clients.
IETF-Radius-Framed-IP-Netmask	Y	Y	Y	String	Single	An IP address mask. The setting applies to VPN remote access IPsec and SSL VPN clients.
IETF-Radius-Idle-Timeout	Y	Y	Y	Integer	Single	Seconds
IETF-Radius-Service-Type	Y	Y	Y	Integer	Single	1 = Login 2 = Framed 5 = Remote access 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 or Active Directory

Table B-2 ASA Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
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
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	String
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; The default is 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

Table B-2 ASA Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
IPsec-User-Group-Lock	Y			Boolean	Single	0 = Disabled 1 = Enabled
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				Integer	Single	For usernames, 0 - 15
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	—

Table B-2 ASA 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 (IKEv1) 8 = L2TP/IPSec 16 = WebVPN 32 = SVC 64 = IPsec (IKEv2) 8 and 4 are mutually exclusive (0 - 11, 16 - 27, 32 - 43, 48 - 59 are legal values).
Use-Client-Address	Y			Boolean	Single	0 = Disabled 1 = Enabled
User-Auth-Server-Name	Y			String	Single	IP address or hostname

Table B-2 ASA Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
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
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 Versions 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://www.example.com
WebVPN-Macro-Substitution-Value1	Y	Y		String	Single	See the *SSL VPN Deployment Guide* for examples at the following 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 the *SSL VPN Deployment Guide* for examples at the following 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

Table B-2 ASA Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
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 Number 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 B-3 describes the syntax rules.

Table B-3 AV-Pair Attribute Syntax Rules

Field	Description
Action	Action to perform if the rule matches a deny or a permit.
Destination	Network or host that receives the packet. Specify it as an IP address, a hostname, or the **any** keyword. 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 of 0 - 65535.
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.
Protocol	Number or name of an IP protocol. Either an integer in the range of 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 **any** keyword. If using an IP address, the source wildcard mask must follow. This field does not apply to Clientless SSL VPN because the ASA has the role of the source or proxy.
Source Wildcard Mask	The wildcard mask that applies to the source address. This field does not apply to Clientless SSL VPN because the ASA has the role of the source or proxy.

Cisco AV Pairs ACL Examples

Table B-4 shows examples of Cisco AV pairs and describes the permit or deny actions that result.

Note Each ACL # in inacl# must be unique. However, they do not need to be sequential (for example, 1, 2, 3, 4). That is, they could be 5, 45, 135.

Table B-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 a 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 a full tunnel IPsec or SSL VPN client.

`webvpn:inacl#1=permit url http://www.example.com`
`webvpn:inacl#2=deny url smtp://server`
`webvpn:inacl#3=permit url cifs://server/share`
	Allows clientlessSSL VPN traffic to the URL specified, denies SMTP traffic to a specific server, and allows file share access (CIFS) to the specified server.

Table B-4 Examples of Cisco AV Pairs and Their Permitting or Denying Action (continued)

Cisco AV Pair Example	Permitting or Denying Action
`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 access and permits SSH access on non-default ports 2323 and 2222, respectively, or other application traffic flows using these ports for clientless SSL VPN.

`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 clientless SSL VPN SSH access to default port 22 and denies Telnet access to port 23, respectively. This example assumes that you are using Telnet or SSH Java plug-ins enforced by these ACLs.

URL Types Supported in ACLs

The URL may be a partial URL, contain wildcards for the server, or include a port.

The following URL types are supported.

any All URLs	https://	post://	ssh://
cifs://	ica://	rdp://	telnet://
citrix://	imap4://	rdp2://	vnc://
citrixs://	ftp://	smart-tunnel://	
http://	pop3://	smtp://	

Note The URLs listed in this table appear in CLI or ASDM menus based on whether or not the associated plug-in 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 do not specify the source because the ASA is the source.

Table B-5 lists the tokens for the Cisco-AV-pair attribute:

Table B-5 ASA-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)

Table B-5 ASA-Supported Tokens (continued)

Token	Syntax Field	Description
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 occurs, 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 Examples

This section presents example procedures for configuring authentication and authorization on the ASA using the Microsoft Active Directory server. It includes the following topics:

- User-Based Attributes Policy Enforcement, page B-16
- Placing LDAP Users in a Specific Group Policy, page B-17
- Enforcing Static IP Address Assignment for AnyConnect Tunnels, page B-19
- Enforcing Dial-in Allow or Deny Access, page B-22
- Enforcing Logon Hours and Time-of-Day Rules, page B-24

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 the following URL:

  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 the following URL:

  http://www.cisco.com/en/US/partner/products/ps6120/products_configuration_example09186a00808d1a7c.shtml

User-Based Attributes Policy Enforcement

You can map any standard LDAP attribute to a well-known Vendor-Specific Attribute (VSA) as well as map one or more LDAP attribute(s) to one or more Cisco LDAP attributes.

The following example shows how to configure the ASA to enforce a simple banner for a user configured on an AD LDAP server. On the server, use the Office field in the General tab to enter the banner text. This field uses the attribute named physicalDeliveryOfficeName. On the ASA, create an attribute map that maps physicalDeliveryOfficeName to the Cisco attribute Banner1. During authentication, the ASA retrieves the value of physicalDeliveryOfficeName from the server, maps the value to the Cisco attribute Banner1, and displays the banner to the user.

This example applies to any connection type, including the IPsec VPN client, AnyConnect SSL VPN client, or clientless SSL VPN. In the example, User1 connects through a clientless SSL VPN connection.

To configure the attributes for a user on the AD or LDAP Server, perform the following steps:

Step 1 Right-click a user.

The Properties dialog box appears (see Figure B-3).

Step 2 Click the **General** tab and enter banner text in the Office field, which uses the AD/LDAP attribute physicalDeliveryOfficeName.

Figure B-3 LDAP User Configuration

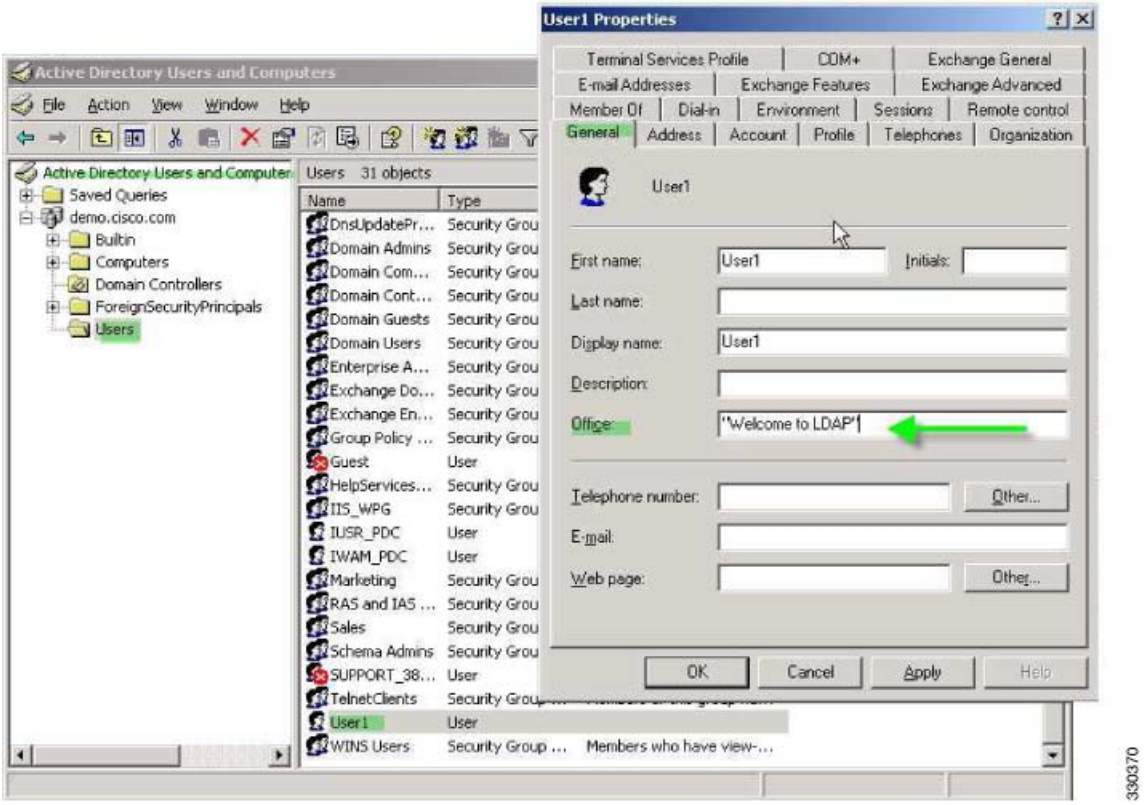

Step 3 Create an LDAP attribute map on the ASA.

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 4 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host 10.1.1.2 in the AAA server group MS_LDAP, and associates the attribute map Banner that you created in Step 3:

```
hostname(config)# aaa-server MS_LDAP host 10.1.1.2
hostname(config-aaa-server-host)# ldap-attribute-map Banner
```

Step 5 Test the banner enforcement.

The following example shows a clientless SSL connection and the banner enforced through the attribute map after the user authenticates (see Figure B-4).

Figure B-4 Banner Displayed

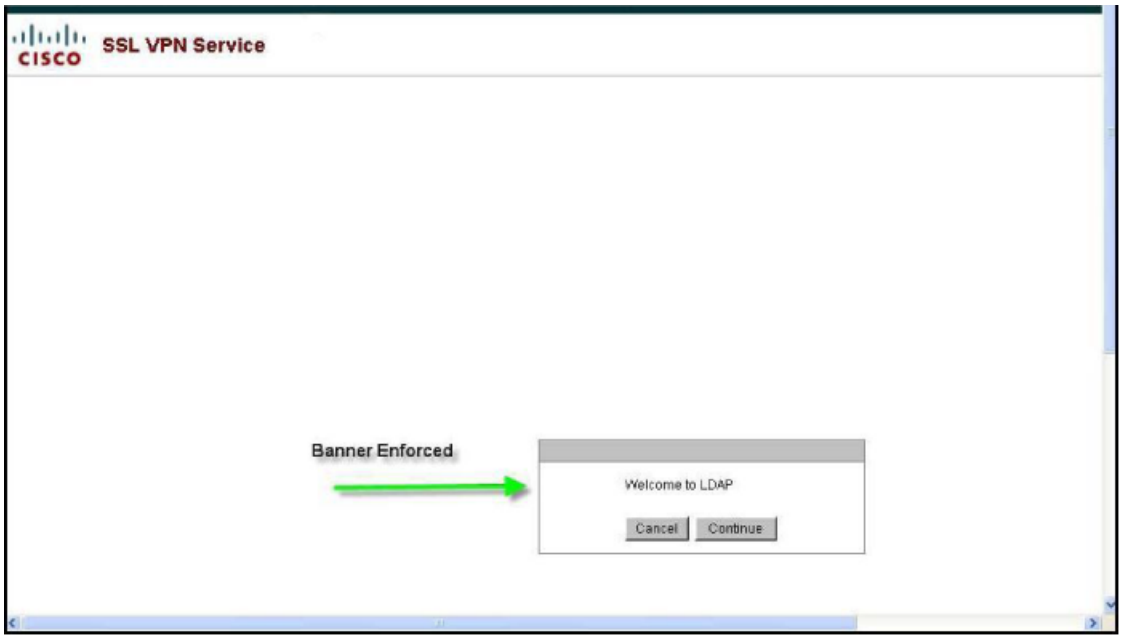

Placing LDAP Users in a Specific Group Policy

The following example shows how to authenticate User1 on the AD LDAP server to a specific group policy on the ASA. On the server, use the Department field of the Organization tab to enter the name of the group policy. Then create an attribute map and map Department to the Cisco attribute IETF-Radius-Class. During authentication, the ASA retrieves the value of Department from the server, maps the value to the IETF-Radius-Class, and places User1 in the group policy.

This example applies to any connection type, including the IPsec VPN client, AnyConnect SSL VPN client, or clientless SSL VPN. In this example, User1 is connecting through a clientless SSL VPN connection.

To configure the attributes for the user on the AD LDAP server, perform the following steps:

Step 1 Right-click the user.

The Properties dialog box appears (see Figure B-5).

Step 2 Click the **Organization** tab and enter **Group-Policy-1** in the Department field.

Figure B-5 AD/LDAP Department Attribute

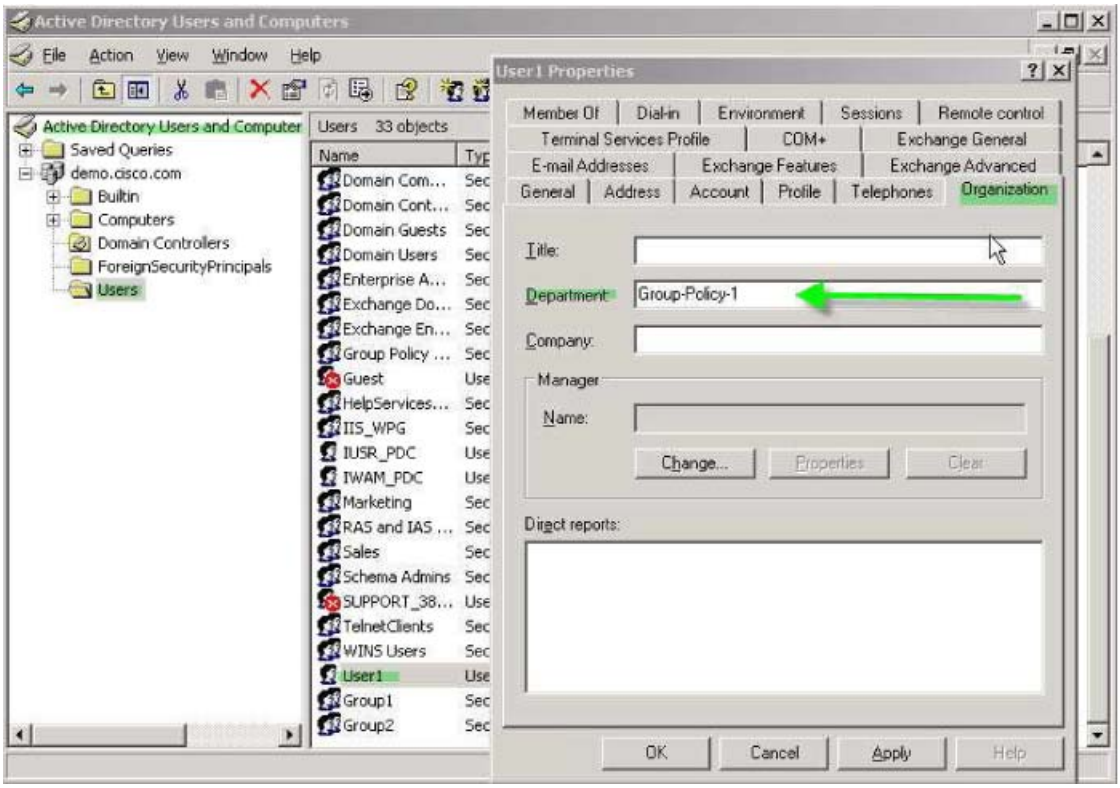

Step 3 Define an attribute map for the LDAP configuration shown in Step 1.

The following example shows how to map the AD attribute Department to the Cisco attribute IETF-Radius-Class.

```
hostname(config)# ldap attribute-map group_policy
hostname(config-ldap-attribute-map)# map-name Department IETF-Radius-Class
```

Step 4 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host 10.1.1.2 in the AAA server group MS_LDAP, and associates the attribute map group_policy that you created in Step 3:

```
hostname(config)# aaa-server MS_LDAP host 10.1.1.2
hostname(config-aaa-server-host)# ldap-attribute-map group_policy
```

Step 5 Add the new group-policy on the ASA and configure the required policy attributes that will be assigned to the user. The following example creates 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 6 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).

Step 7 Monitor the communication between the ASA and the server by enabling the **debug ldap 255** command from privileged EXEC mode. The following is sample output from this command, which has been edited to provide the key messages:

```
[29] Authentication successful for user1 to 10.1.1.2
[29] Retrieving user attributes from server 10.1.1.2
[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 example, configure the AnyConnect client user Web1 to receive a static IP address. then 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. Create an attribute map that maps this attribute to the Cisco attribute IETF-Radius-Framed-IP-Address.

During authentication, the ASA 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.

The following example applies to full-tunnel clients, including the IPsec client and the SSL VPN clients (AnyConnect client 2.x and the SSL VPN client).

To configure the user attributes on the AD /LDAP server, perform the following steps:

Step 1 Right-click the username.

The Properties dialog box appears (see Figure B-6).

Step 2 Click the **Dialin** tab, check the **Assign Static IP Address** check box, and enter an IP address of 10.1.1.2.

Figure B-6 Assign Static IP Address

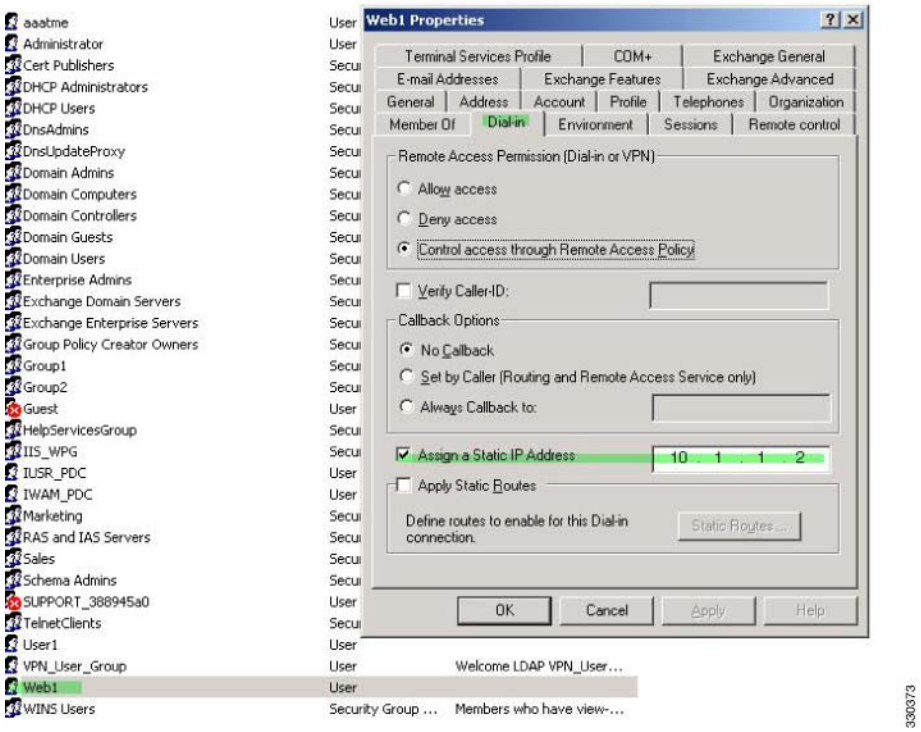

Step 3 Create an attribute map for the LDAP configuration shown in Step 1.

The following example shows how to map the AD attribute msRADIUSFramedIPAddress used by the Static Address field to the Cisco attribute IETF-Radius-Framed-IP-Address:

```
hostname(config)# ldap attribute-map static_address
hostname(config-ldap-attribute-map)# map-name msRADIUSFramedIPAddress
IETF-Radius-Framed-IP-Address
```

Step 4 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host 10.1.1.2, in the AAA server group MS_LDAP, and associates the attribute map static_address that you created in Step 3:

```
hostname(config)# aaa-server MS_LDAP host 10.1.1.2
hostname(config-aaa-server-host)# ldap-attribute-map static_address
```

Step 5 Verify that the **vpn-address-assignment** command is configured to specify AAA by viewing this part of the configuration with the **show run all vpn-addr-assign** command:

```
hostname(config)# show run all vpn-addr-assign
vpn-addr-assign aaa    << Make sure this is configured >>
no vpn-addr-assign dhcp
vpn-addr-assign local
hostname(config)#
```

Step 6 Establish a connection to the ASA with the AnyConnect client. Observe the following:

- The banner is received in the same sequence as a clientless connection (see Figure B-7).
- The user receives the IP address configured on the server and mapped to the ASA (see Figure B-8).

Figure B-7 Verify the Banner for the AnyConnect Session

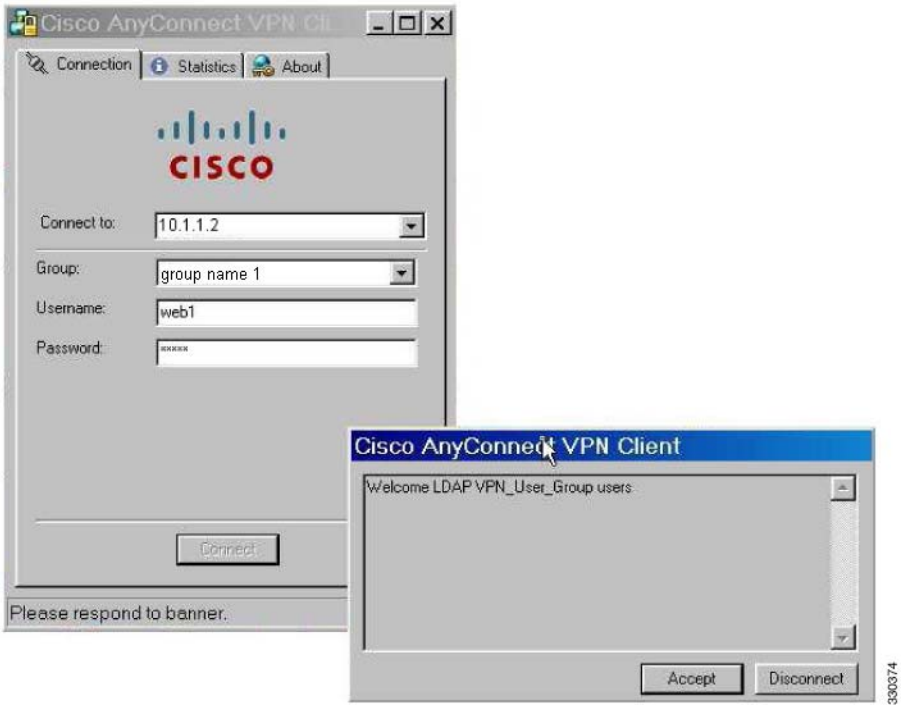

Figure B-8 AnyConnect Session Established

Step 7 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  : 10.1.1.2               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 : Group1_TunnelGroup
Login Time   : 11:13:05 UTC Tue Aug 28 2007
Duration     : 0h:01m:48s
NAC Result   : Unknown
VLAN Mapping : N/A                     VLAN         : none
```

Enforcing Dial-in Allow or Deny Access

The following example creates an LDAP attribute map that specifies the tunneling protocols allowed by the user. You map the allow access and deny access settings on the Dialin tab to the Cisco attribute Tunneling-Protocol, which supports the bitmap values shown in Table B-6:

Table B-6 Bitmap Values for Cisco Tunneling-Protocol Attribute

Value	Tunneling Protocol
1	PPTP
2	L2TP
4[1]	IPsec (IKEv1)
8[2]	L2TP/IPsec
16	Clientless SSL
32	SSL client—AnyConnect or SSL VPN client
64	IPsec (IKEv2)

1. IPsec and L2TP over IPsec are not supported simultaneously. Therefore, the values 4 and 8 are mutually exclusive.
2. See note 1.

Use this attribute to create an Allow Access (TRUE) or a Deny Access (FALSE) condition for the protocols and enforce the method for which the user is allowed access.

For this simplified example, by mapping the tunnel protocol IPsec/IKEv1 (4), you can create an allow (true) condition for the Cisco VPN client. You also map WebVPN (16) and SVC/AC (32), which are mapped as a value of 48 (16+32) and create a deny (false) condition. This allows the user to connect to the ASA using IPsec, but any attempt to connect using clientless SSL or the AnyConnect client is denied.

Another example of enforcing dial-in allow access or deny access is available in the Tech Note *ASA/PIX: Mapping VPN Clients to VPN Group Policies Through LDAP Configuration Example* at the following URL:

http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a008089149d.shtml

To configure the user attributes on the AD/LDAP server, perform the following steps:

Step 1 Right-click the user.

The Properties dialog box appears.

Step 2 Click the **Dial-in** tab, then click the **Allow Access** radio button (Figure B-9).

Figure B-9 AD/LDAP User1 - Allow Access

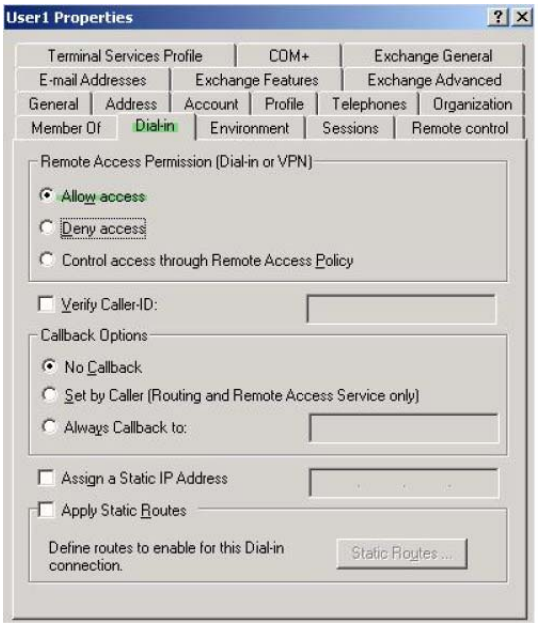

Note If you select the Control access through the Remote Access Policy option, then a value is not returned from the server, and the permissions that are enforced are based on the internal group policy settings of the ASA.

Step 3 Create an attribute map to allow both an IPsec and AnyConnect connection, but deny a clientless SSL connection.

The following example shows how to 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:

```
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 4 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host 10.1.1.2, 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 10.1.1.2
hostname(config-aaa-server-host)# ldap-attribute-map tunneling_protocols
```

Step 5 Verify that the attribute map works as configured.

Step 6 Try 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 the attribute map (see Figure B-10 and Figure B-11).

Figure B-10 Login Denied Message for Clientless User

Figure B-11 Login Denied Message for AnyConnect Client User

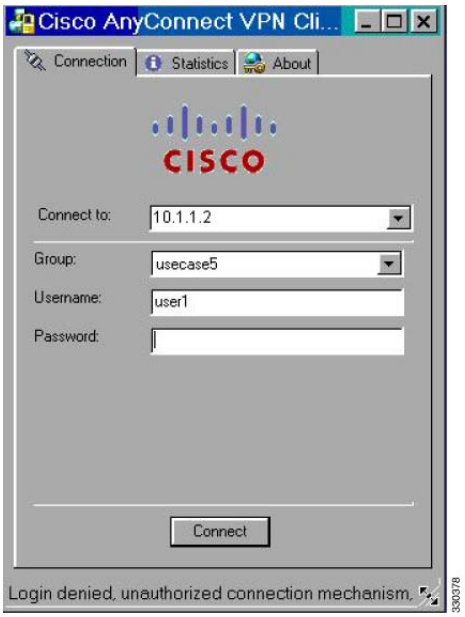

Enforcing Logon Hours and Time-of-Day Rules

The following example shows how to configure and enforce the hours that a clientless SSL user (such as a business partner) is allowed to access the network.

On the AD server, use the Office field to enter the name of the partner, which uses the physicalDeliveryOfficeName attribute. Then we create an attribute map on the ASA to map that attribute to the Cisco attribute Access-Hours. During authentication, the ASA retrieves the value of physicalDeliveryOfficeName and maps it to Access-Hours.

To configure the user attributes on the AD /LDAP server, perform the following steps:

Step 1 Select the user, and right-click **Properties**.

The Properties dialog box appears (see Figure B-12).

Step 2 Click the **General** tab.

Figure B-12 Active Directory Properties Dialog Box

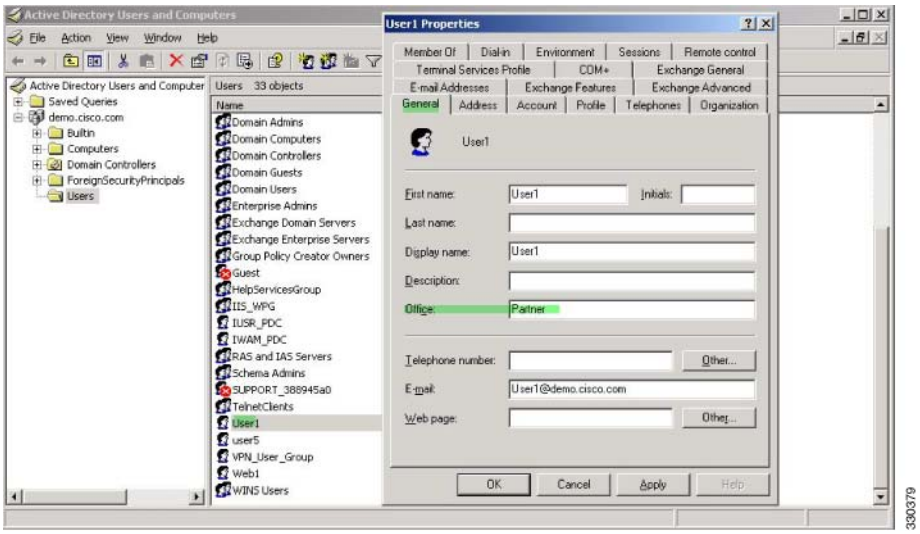

Step 3 Create an attribute map.

The following example shows how to create the attribute map access_hours and map the AD attribute physicalDeliveryOfficeName used by the Office field to the Cisco attribute Access-Hours.

```
hostname(config)# ldap attribute-map access_hours
hostname(config-ldap-attribute-map)# map-name physicalDeliveryOfficeName Access-Hours
```

Step 4 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host 10.1.1.2, in the AAA server group MS_LDAP, and associates the attribute map access_hours that you created in Step 3:

```
hostname(config)# aaa-server MS_LDAP host 10.1.1.2
hostname(config-aaa-server-host)# ldap-attribute-map access_hours
```

Step 5 Configure time ranges for each value allowed on the server.

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 B-26
- ASA RADIUS Authorization Attributes, page B-26
- ASA IETF RADIUS Authorization Attributes, page B-36
- RADIUS Accounting Disconnect Reason Codes, page B-36

Reviewing the RADIUS Configuration Procedure

This section describes the RADIUS configuration steps required to support authentication and authorization of ASA users.

To set up the RADIUS server to interoperate with the ASA, perform the following steps:

Step 1 Load the ASA 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.
- For RADIUS servers from other vendors (for example, Microsoft Internet Authentication Service): you must manually define each ASA attribute. To define an attribute, use the attribute name or number, type, value, and vendor code (3076). For a list of ASA RADIUS authorization attributes and values, see Table B-7.

Step 2 Set up the users or groups with the permissions and attributes to send during IPsec or SSL tunnel establishment.

ASA 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. These attributes have vendor ID 3076.

Table B-7 lists the ASA 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 ASAs enforce the RADIUS attributes based on attribute numeric ID, not attribute name. LDAP attributes are enforced by their name, not by the ID.

All attributes listed in Table B-7 are downstream attributes that are sent from the RADIUS server to the ASA except for the following attribute numbers: 146, 150, 151, and 152. These attribute numbers are upstream attributes that are sent from the ASA to the RADIUS server. RADIUS attributes 146 and 150 are sent from the ASA to the RADIUS server for authentication and authorization requests. All four previously listed attributes are sent from the ASA to the RADIUS server for accounting start, interim-update, and stop requests. Upstream RADIUS attributes 146, 150, 151, and 152 were introduced in ASA version 8.4.3.
Cisco ACS 5.x and Cisco ISE do not support IPv6 framed IP addresses for IP address assignment using RADIUS authentication in ASA Version 9.0.

Table B-7 ASA Supported RADIUS Attributes and Values

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
Access-Hours	Y	1	String	Single	Name of the time range, for example, Business-hours
Access-List-Inbound	Y	86	String	Single	ACL ID
Access-List-Outbound	Y	87	String	Single	ACL ID
Address-Pools	Y	217	String	Single	Name of IP local pool
Allow-Network-Extension-Mode	Y	64	Boolean	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle-Timeout	Y	50	Integer	Single	1-35791394 minutes
Authorization-DN-Field	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
Authorization-Required		66	Integer	Single	0 = No 1 = Yes
Authorization-Type	Y	65	Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Banner1	Y	15	String	Single	Banner string to display for Cisco VPN remote access sessions: IPsec IKEv1, AnyConnect SSL-TLS/DTLS/IKEv2, and Clientless SSL
Banner2	Y	36	String	Single	Banner string to display for Cisco VPN remote access sessions: IPsec IKEv1, AnyConnect SSL-TLS/DTLS/IKEv2, and Clientless SSL. The Banner2 string is concatenated to the Banner1 string , if configured.
Cisco-IP-Phone-Bypass	Y	51	Integer	Single	0 = Disabled 1 = Enabled
Cisco-LEAP-Bypass	Y	75	Integer	Single	0 = Disabled 1 = Enabled
Client Type	Y	150	Integer	Single	1 = Cisco VPN Client (IKEv1) 2 = AnyConnect Client SSL VPN 3 = Clientless SSL VPN 4 = Cut-Through-Proxy 5 = L2TP/IPsec SSL VPN 6 = AnyConnect Client IPsec VPN (IKEv2)
Client-Type-Version-Limiting	Y	77	String	Single	IPsec VPN version number string
DHCP-Network-Scope	Y	61	String	Single	IP Address
Extended-Authentication-On-Rekey	Y	122	Integer	Single	0 = Disabled 1 = Enabled

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
Group-Policy	Y	25	String	Single	Sets the group policy for the remote access VPN session. For Versions 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		83	Integer	Single	0 = None 1 = Local
IE-Proxy-Exception-List		82	String	Single	New line (\n) separated list of DNS domains
IE-Proxy-PAC-URL	Y	133	String	Single	PAC Address String
IE-Proxy-Server		80	String	Single	IP address
IE-Proxy-Server-Policy		81	Integer	Single	1 = No Modify 2 = No Proxy 3 = Auto detect 4 = Use Concentrator Setting
IKE-KeepAlive-Confidence-Interval	Y	68	Integer	Single	10 - 300 seconds
IKE-Keepalive-Retry-Interval	Y	84	Integer	Single	2 - 10 seconds
IKE-Keep-Alives	Y	41	Boolean	Single	0 = Disabled 1 = Enabled
Intercept-DHCP-Configure-Msg	Y	62	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Allow-Passwd-Store	Y	16	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Authentication		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
IPsec-Auth-On-Rekey	Y	42	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Backup-Server-List	Y	60	String	Single	Server Addresses (space delimited)
IPsec-Backup-Servers	Y	59	String	Single	1 = Use Client-Configured list 2 = Disable and clear client list 3 = Use Backup Server list

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
IPsec-Client-Firewall-Filter-Name		57	String	Single	Specifies the name of the filter to be pushed to the client as firewall policy
IPsec-Client-Firewall-Filter-Optional	Y	58	Integer	Single	0 = Required 1 = Optional
IPsec-Default-Domain	Y	28	String	Single	Specifies the single default domain name to send to the client (1-255 characters).
IPsec-IKE-Peer-ID-Check	Y	40	Integer	Single	1 = Required 2 = If supported by peer certificate 3 = Do not check
IPsec-IP-Compression	Y	39	Integer	Single	0 = Disabled 1 = Enabled
IPsec-Mode-Config	Y	31	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP-Port	Y	35	Integer	Single	4001 - 49151. The default is10000.
IPsec-Required-Client-Firewall-Capability	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-Sec-Association		12	String	Single	Name of the security association
IPsec-Split-DNS-Names	Y	29	String	Single	Specifies the list of secondary domain names to send to the client (1-255 characters).
IPsec-Split-Tunneling-Policy	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Local LAN permitted
IPsec-Split-Tunnel-List	Y	27	String	Single	Specifies the name of the network/ACL that describes the split tunnel inclusion list.
IPsec-Tunnel-Type	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPsec-User-Group-Lock		33	Boolean	Single	0 = Disabled 1 = Enabled
IPv6-Address-Pools	Y	218	String	Single	Name of IP local pool-IPv6
IPv6-VPN-Filter	Y	219	String	Single	ACL value

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
L2TP-Encryption		21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Req 15= 40/128-Encr/Stateless-Req
L2TP-MPPC-Compression		38	Integer	Single	0 = Disabled 1 = Enabled
Member-Of	Y	145	String	Single	Comma-delimited string, for example: `Engineering, Sales` An administrative attribute that can be used in dynamic access policies. It does not set a group policy.
MS-Client-Subnet-Mask	Y	63	Boolean	Single	An IP address
NAC-Default-ACL		92	String		ACL
NAC-Enable		89	Integer	Single	0 = No 1 = Yes
NAC-Revalidation-Timer		91	Integer	Single	300 - 86400 seconds
NAC-Settings	Y	141	String	Single	Name of the NAC policy
NAC-Status-Query-Timer		90	Integer	Single	30 - 1800 seconds
Perfect-Forward-Secrecy-Enable	Y	88	Boolean	Single	0 = No 1 = Yes
PPTP-Encryption		20	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Required 15= 40/128-Encr/Stateless-Req
PPTP-MPPC-Compression		37	Integer	Single	0 = Disabled 1 = Enabled
Primary-DNS	Y	5	String	Single	An IP address
Primary-WINS	Y	7	String	Single	An IP address
Privilege-Level	Y	220	Integer	Single	An integer between 0 and 15.
Required-Client- Firewall-Vendor-Code	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)

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
Required-Client-Firewall-Description	Y	47	String	Single	String
Required-Client-Firewall-Product-Code	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-Individual-User-Auth	Y	49	Integer	Single	0 = Disabled 1 = Enabled
Require-HW-Client-Auth	Y	48	Boolean	Single	0 = Disabled 1 = Enabled
Secondary-DNS	Y	6	String	Single	An IP address
Secondary-WINS	Y	8	String	Single	An IP address
SEP-Card-Assignment		9	Integer	Single	Not used
Session Subtype	Y	152	Integer	Single	0 = None 1 = Clientless 2 = Client 3 = Client Only Session Subtype applies only when the Session Type (151) attribute has the following values: 1, 2, 3, and 4.
Session Type	Y	151	Integer	Single	0 = None 1 = AnyConnect Client SSL VPN 2 = AnyConnect Client IPSec VPN (IKEv2) 3 = Clientless SSL VPN 4 = Clientless Email Proxy 5 = Cisco VPN Client (IKEv1) 6 = IKEv1 LAN-LAN 7 = IKEv2 LAN-LAN 8 = VPN Load Balancing
Simultaneous-Logins	Y	2	Integer	Single	0 - 2147483647
Smart-Tunnel	Y	136	String	Single	Name of a Smart Tunnel

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
Strip-Realm	Y	135	Boolean	Single	0 = Disabled 1 = Enabled
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
SVC-DPD-Interval-Client	Y	108	Integer	Single	0 = Off 5 - 3600 seconds
SVC-DPD-Interval-Gateway	Y	109	Integer	Single	0 = Off) 5 - 3600 seconds
SVC-DTLS	Y	123	Integer	Single	0 = False 1 = True
SVC-Keepalive	Y	107	Integer	Single	0 = Off 15 - 600 seconds
SVC-Modules	Y	127	String	Single	String (name of a module)
SVC-MTU	Y	125	Integer	Single	MTU value 256 - 1406 in bytes
SVC-Profiles	Y	128	String	Single	String (name of a profile)
SVC-Rekey-Time	Y	110	Integer	Single	0 = Disabled 1- 10080 minutes
Tunnel Group Name	Y	146	String	Single	1 - 253 characters
Tunnel-Group-Lock	Y	85	String	Single	Name of the tunnel group or “none”
Tunneling-Protocols	Y	11	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec (IKEv1) 8 = L2TP/IPSec 16 = WebVPN 32 = SVC 64 = IPsec (IKEv2) 8 and 4 are mutually exclusive (0 - 11, 16 - 27, 32 - 43, 48 - 59 are legal values).

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
Use-Client-Address		17	Boolean	Single	0 = Disabled 1 = Enabled
VLAN	Y	140	Integer	Single	0 - 4094
WebVPN-Access-List	Y	73	String	Single	Access-List name
WebVPN ACL	Y	73	String	Single	Name of a WebVPN ACL on the device
WebVPN-ActiveX-Relay	Y	137	Integer	Single	0 = Disabled Otherwise = Enabled
WebVPN-Apply-ACL	Y	102	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Auto-HTTP-Signon	Y	124	String	Single	Reserved
WebVPN-Citrix-Metaframe-Enable	Y	101	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Content-Filter-Parameters	Y	69	Integer	Single	1 = Java ActiveX 2 = Java Script 4 = Image 8 = Cookies in images
WebVPN-Customization	Y	113	String	Single	Name of the customization
WebVPN-Default-Homepage	Y	76	String	Single	A URL such as http://example-example.com
WebVPN-Deny-Message	Y	116	String	Single	Valid string (up to 500 characters)
WebVPN-Download_Max-Size	Y	157	Integer	Single	0x7fffffff
WebVPN-File-Access-Enable	Y	94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Group-based-HTTP/HTTPS-Proxy -Exception-List	Y	78	String	Single	Comma-separated DNS/IP with an optional wildcard (*) (for example *.cisco.com, 192.168.1.*, wwwin.cisco.com)
WebVPN-Hidden-Shares	Y	126	Integer	Single	0 = None 1 = Visible
WebVPN-Home-Page-Use-Smart-Tunnel	Y	228	Boolean	Single	Enabled if clientless home page is to be rendered through Smart Tunnel.
WebVPN-HTML-Filter	Y	69	Bitmap	Single	1 = Java ActiveX 2 = Scripts 4 = Image 8 = Cookies
WebVPN-HTTP-Compression	Y	120	Integer	Single	0 = Off 1 = Deflate Compression

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
WebVPN-HTTP-Proxy-IP-Address	Y	74	String	Single	Comma-separated DNS/IP:port, with http= or https= prefix (for example http=10.10.10.10:80, https=11.11.11.11:443)
WebVPN-Idle-Timeout-Alert-Interval	Y	148	Integer	Single	0 (Disabled) - 30
WebVPN-Keepalive-Ignore	Y	121	Integer	Single	0-900
WebVPN-Macro-Substitution	Y	223	String	Single	Unbounded. For examples, see the *SSL VPN Deployment Guide* at the following URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x
WebVPN-Macro-Substitution	Y	224	String	Single	Unbounded. For examples, see the *SSL VPN Deployment Guide* at the following URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x
WebVPN-Port-Forwarding-Enable	Y	97	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-List	Y	72	String	Single	Port forwarding list name
WebVPN-Port-Forwarding-Name	Y	79	String	Single	String name (example, “Corporate-Apps”). This text replaces the default string, “Application Access,” on the clientless portal home page.
WebVPN-Post-Max-Size	Y	159	Integer	Single	0x7fffffff
WebVPN-Session-Timeout-Alert-Interval	Y	149	Integer	Single	0 (Disabled) - 30
WebVPN Smart-Card-Removal-Disconnect	Y	225	Boolean	Single	0 = Disabled 1 = Enabled
WebVPN-Smart-Tunnel	Y	136	String	Single	Name of a smart tunnel
WebVPN-Smart-Tunnel-Auto-Sign-On	Y	139	String	Single	Name of a Smart Tunnel auto sign-on list appended by the domain name
WebVPN-Smart-Tunnel-Auto-Start	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = Auto Start

Table B-7 ASA Supported RADIUS Attributes and Values (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
WebVPN-Smart-Tunnel-Tunnel-Policy	Y	227	String	Single	One of "e networkname," "i networkname," or "a," where networkname is the name of a smart tunnel network list, e indicates the tunnel excluded, i indicates the tunnel specified, and a indicates all tunnels.
WebVPN-SSL-VPN-Client-Enable	Y	103	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Keep-Installation	Y	105	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Required	Y	104	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSO-Server-Name	Y	114	String	Single	Valid string
WebVPN-Storage-Key	Y	162	String	Single	
WebVPN-Storage-Objects	Y	161	String	Single	
WebVPN-SVC-Keepalive-Frequency	Y	107	Integer	Single	15-600 seconds, 0=Off
WebVPN-SVC-Client-DPD-Frequency	Y	108	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-DTLS-Enable	Y	123	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SVC-DTLS-MTU	Y	125	Integer	Single	MTU value is from 256-1406 bytes.
WebVPN-SVC-Gateway-DPD-Frequency	Y	109	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-Rekey-Time	Y	110	Integer	Single	4-10080 minutes, 0=Off
WebVPN-SVC-Rekey-Method	Y	111	Integer	Single	0 (Off), 1 (SSL), 2 (New Tunnel)
WebVPN-SVC-Compression	Y	112	Integer	Single	0 (Off), 1 (Deflate Compression)
WebVPN-UNIX-Group-ID (GID)	Y	222	Integer	Single	Valid UNIX group IDs
WebVPN-UNIX-User-ID (UIDs)	Y	221	Integer	Single	Valid UNIX user IDs
WebVPN-Upload-Max-Size	Y	158	Integer	Single	0x7fffffff
WebVPN-URL-Entry-Enable	Y	93	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-List	Y	71	String	Single	URL list name
WebVPN-User-Storage	Y	160	String	Single	
WebVPN-VDI	Y	163	String	Single	List of settings

ASA IETF RADIUS Authorization Attributes

Table B-8 lists the supported IETF RADIUS attributes.

Table B-8 ASA Supported IETF RADIUS Attributes and Values

Attribute Name	VPN 3000	ASA	PIX	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
IETF-Radius-Class	Y	Y	Y	25		Single	For Versions 8.2.x and later, we recommend that you use the Group-Policy attribute (VSA 3076, #25) as described in Table B-7: • *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 ASA, which 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
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

RADIUS Accounting Disconnect Reason Codes

These codes are returned if the ASA encounters a disconnect when sending packets:

Table B-9

Disconnect Reason Code
ACCT_DISC_USER_REQ = 1
ACCT_DISC_LOST_CARRIER = 2
ACCT_DISC_LOST_SERVICE = 3
ACCT_DISC_IDLE_TIMEOUT = 4
ACCT_DISC_SESS_TIMEOUT = 5

Table B-9

Disconnect Reason Code
ACCT_DISC_ADMIN_RESET = 6
ACCT_DISC_ADMIN_REBOOT = 7
ACCT_DISC_PORT_ERROR = 8
ACCT_DISC_NAS_ERROR = 9
ACCT_DISC_NAS_REQUEST = 10
ACCT_DISC_NAS_REBOOT = 11
ACCT_DISC_PORT_UNNEEDED = 12
ACCT_DISC_PORT_PREEMPTED = 13
ACCT_DISC_PORT_SUSPENDED = 14
ACCT_DISC_SERV_UNAVAIL = 15
ACCT_DISC_CALLBACK = 16
ACCT_DISC_USER_ERROR = 17
ACCT_DISC_HOST_REQUEST = 18
ACCT_DISC_ADMIN_SHUTDOWN = 19
ACCT_DISC_SA_EXPIRED = 21
ACCT_DISC_MAX_REASONS = 22

Configuring an External TACACS+ Server

The ASA 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 that you have enabled AAA services on the NAS.

Table B-10 lists supported TACACS+ authorization response attributes for cut-through-proxy connections. Table B-11 lists supported TACACS+ accounting attributes.

Table B-10 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 B-11 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).
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 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.

Numerics

3DES See DES.

A

AAA Authentication, authorization, and accounting. See also TACACS+ and RADIUS.

ABR Area Border Router. In OSPF, a router with interfaces in multiple areas.

ACE access control entry. Information entered into the configuration that lets you specify what type of traffic to permit or deny on an interface. By default, traffic that is not explicitly permitted is denied.

Access Modes The ASA CLI uses several command modes. The commands available in each mode vary. See also user EXEC mode, privileged EXEC mode, global configuration mode, command-specific configuration mode.

ACL access control list. A collection of ACEs. An ACL lets you specify what type of traffic to allow on an interface. By default, traffic that is not explicitly permitted is denied. ACLs are usually applied to the interface which is the source of inbound traffic. See also rule, outbound ACL.

ActiveX A set of object-oriented programming technologies and tools used to create mobile or portable programs. An ActiveX program is roughly equivalent to a Java applet.

Address Resolution Protocol See ARP.

address translation The translation of a network address and/or port to another network address/or port. See also IP address, interface PAT, NAT, PAT, Static PAT, xlate.

AES Advanced Encryption Standard. A symmetric block cipher that can encrypt and decrypt information. The AES algorithm is capable of using cryptographic keys of 128, 192 and 256 bits to encrypt and decrypt data in blocks of 128 bits. See also DES.

AH Authentication Header. An IP protocol (type 51) that can ensure data integrity, authentication, and replay detection. AH is embedded in the data to be protected (a full IP datagram, for example). AH can be used either by itself or with ESP. AH is an older IPsec protocol that is less important in most networks than ESP. AH provides authentication services but does not provide encryption services. It is provided to ensure compatibility with IPsec peers that do not support ESP, which provides both authentication and encryption. See also encryption and VPN. Refer to the RFC 2402.

AIP Advanced Inspection and Prevention. For example, the AIP SSM or AIP SSC, which runs IPS software.

A record address "A" stands for address, and refers to name-to-address mapped records in DNS.

APCF Application Profile Customization Framework. Lets the security appliance handle nonstandard applications so that they render correctly over a clientless SSL VPN connection.

ARP Address Resolution Protocol. A low-level TCP/IP protocol that maps a hardware address, or MAC address, to an IP address. An example hardware address is 00:00:a6:00:01:ba. The first three groups of characters (00:00:a6) identify the manufacturer; the rest of the characters (00:01:ba) identify the system card. ARP is defined in RFC 826.

ASA Adaptive Security Algorithm. Used by the ASA to perform inspections. ASA allows one-way (inside to outside) connections without an explicit configuration for each internal system and application. See also inspection engine.

ASA adaptive ASA.

ASDM Adaptive Security Device Manager. An application for managing and configuring a single ASA.

asymmetric encryption Also called public key systems, asymmetric encryption allows anyone to obtain access to the public key of anyone else. Once the public key is accessed, you can send an encrypted message to that person using the public key. See also encryption, public key.

authentication Cryptographic protocols and services that verify the identity of users and the integrity of data. One of the functions of the IPsec framework. Authentication establishes the integrity of the datastream and ensures that it is not tampered with in transit. It also provides confirmation about the origin of the datastream. See also AAA, encryption, and VPN.

Auto Applet Download Automatically downloads the clientless SSL VPN port-forwarding applet when the user first logs in to clientless SSL VPN.

auto-signon This command provides a single sign-on method for clientless SSL VPN users. It passes the clientless SSL VPN login credentials (username and password) to internal servers for authentication using NTLM authentication, basic authentication, or both.

B

backup server IPsec backup servers let a VPN client connect to the central site when the primary security appliance is unavailable.

BGP Border Gateway Protocol. BGP performs interdomain routing in TCP/IP networks. BGP is an Exterior Gateway Protocol, which means that it performs routing between multiple autonomous systems or domains and exchanges routing and access information with other BGP systems. The ASA does not support BGP. See also EGP.

BLT stream Bandwidth Limited Traffic stream. Stream or flow of packets whose bandwidth is constrained.

BOOTP Bootstrap Protocol. Lets diskless workstations boot over the network as is described in RFC 951 and RFC 1542.

BPDU Bridge Protocol Data Unit. Spanning-Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network. Protocol data unit is the OSI term for packet.

C

CA Certificate Authority, Certification Authority. A third-party entity that is responsible for issuing and revoking certificates. Each device with the public key of the CA can authenticate a device that has a certificate issued by the CA. The term CA also refers to software that provides CA services. See also certificate, CRL, public key, RA.

cache A temporary repository of information accumulated from previous task executions that can be reused, decreasing the time required to perform the tasks. Caching stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content.

CBC Cipher Block Chaining. A cryptographic technique that increases the encryption strength of an algorithm. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.

certificate A signed cryptographic object that contains the identity of a user or device and the public key of the CA that issued the certificate. Certificates have an expiration date and may also be placed on a CRL if known to be compromised. Certificates also establish non-repudiation for IKE negotiation, which means that you can prove to a third party that IKE negotiation was completed with a specific peer.

CHAP Challenge Handshake Authentication Protocol.

CIFS Common Internet File System. It is a platform-independent file sharing system that provides users with network access to files, printers, and other machine resources. Microsoft implemented CIFS for networks of Windows computers, however, open source implementations of CIFS provide file access to servers running other operating systems, such as Linux, UNIX, and Mac OS X.

Citrix An application that virtualizes client-server applications and optimizes web applications.

CLI command-line interface. The primary interface for entering configuration and monitoring commands to the ASA.

client/server computing Distributed computing (processing) network systems in which transaction responsibilities are divided into two parts: client (front end) and server (back end). Also called distributed computing. See also RPC.

Client update Lets you update revisions of clients to which the update applies; provide a URL or IP address from which to get the update; and, in the case of Windows clients, optionally notify users that they should update their VPN client version.

command-specific configuration mode From global configuration mode, some commands enter a command-specific configuration mode. All user EXEC, privileged EXEC, global configuration, and command-specific configuration commands are available in this mode. See also global configuration mode, privileged EXEC mode, user EXEC mode.

compression The process of encoding information using fewer bits or other information-bearing units than an unencoded representation would use. Compression can reduce the size of transferring packets and increase communication performance.

configuration, config, config file A file on the ASA that represents the equivalent of settings, preferences, and properties administered by ASDM or the CLI.

Content Rewriting/Transformation Interprets and modifies applications so that they render correctly over a clientless SSL VPN connection.

cookie A cookie is a object stored by a browser. Cookies contain information, such as user preferences, to persistent storage.

CPU Central Processing Unit. Main processor.

CRC Cyclical Redundancy Check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node.

CRL Certificate Revocation List. A digitally signed message that lists all of the current but revoked certificates listed by a given CA. A CRL is analogous to a book of stolen charge card numbers that allow stores to reject bad credit cards. When certificates are revoked, they are added to a CRL. When you implement authentication using certificates, you can choose to use CRLs or not. Using CRLs lets you easily revoke certificates before they expire, but the CRL is generally only maintained by the CA or an RA. If you are using CRLs and the connection to the CA or RA is not available when authentication is requested, the authentication request will fail. See also CA, certificate, public key, RA.

CRV Call Reference Value. Used by H.225.0 to distinguish call legs signaled between two entities.

cryptography Encryption, authentication, integrity, keys and other services used for secure communication over networks. See also VPN and IPsec.

crypto map A data structure with a unique name and sequence number that is used for configuring VPNs on the ASA. A crypto map selects data flows that need security processing and defines the policy for these flows and the crypto peer that traffic needs to go to. A crypto map is applied to an interface. Crypto maps contain the ACLs, encryption standards, peers, and other parameters necessary to specify security policies for VPNs using IKE and IPsec. See also VPN.

CTIQBE Computer Telephony Interface Quick Buffer Encoding. A protocol used in IP telephony between the Cisco CallManager and CTI TAPI and JTAPI applications. CTIQBE is used by the TAPI/JTAPI protocol inspection module and supports NAT, PAT, and bidirectional NAT. This protocol enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to communicate with Cisco CallManager for call setup and voice traffic across the ASA.

cut-through proxy Enables the ASA to provide faster traffic flow after user authentication. The cut-through proxy challenges a user initially at the application layer. 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.

D

data confidentiality Describes any method that manipulates data so that no attacker can read it. This is commonly achieved by data encryption and keys that are only available to the parties involved in the communication.

data integrity Describes mechanisms that, through the use of encryption based on secret key or public key algorithms, allow the recipient of a piece of protected data to verify that the data has not been modified in transit.

data origin authentication A security service where the receiver can verify that protected data could have originated only from the sender. This service requires a data integrity service plus a key distribution mechanism, where a secret key is shared only between the sender and receiver.

decryption Application of a specific algorithm or cipher to encrypted data so as to render the data comprehensible to those who are authorized to see the information. See also encryption.

DES Data encryption standard. DES was published in 1977 by the National Bureau of Standards and is a secret key encryption scheme based on the Lucifer algorithm from IBM. Cisco uses DES in classic crypto (40-bit and 56-bit key lengths), IPsec crypto (56-bit key), and 3DES (triple DES), which performs encryption three times using a 56-bit key. 3DES is more secure than DES but requires more processing for encryption and decryption. See also AES, ESP.

DHCP Dynamic Host Configuration Protocol. Provides a mechanism for allocating IP addresses to hosts dynamically, so that addresses can be reused when hosts no longer need them and so that mobile computers, such as laptops, receive an IP address applicable to the LAN to which it is connected.

Diffie-Hellman A public key cryptography protocol that allows two parties to establish a shared secret over insecure communications channels. Diffie-Hellman is used within IKE to establish session keys. Diffie-Hellman is a component of Oakley key exchange.

Diffie-Hellman Group 1, Group 2, Group 5, Group 7 Diffie-Hellman refers to a type of public key cryptography using asymmetric encryption based on large prime numbers to establish both Phase 1 and Phase 2 SAs. Group 1 provides a smaller prime number than Group 2 but may be the only version supported by some IPsec peers. Diffe-Hellman Group 5 uses a 1536-bit prime number, is the most secure, and is recommended for use with AES. Group 7 has an elliptical curve field size of 163 bits and is for use with the Movian VPN client, but works with any peer that supports Group 7 (ECC). See also VPN and encryption.

Note The **group 7** command option was deprecated in ASA Version 8.0(4). Attempts to configure group 7 will generate an error message and use group 5 instead.

digital certificate See certificate.

DMZ See interface.

DN Distinguished Name. Global, authoritative name of an entry in the OSI Directory (X.500).

DNS Domain Name System (or Service). An Internet service that translates domain names into IP addresses.

DoS Denial of Service. A type of network attack in which the goal is to render a network service unavailable.

DSL digital subscriber line. Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. DSL is provisioned via modem pairs, with one modem located at a central office and the other at the customer site. Because most DSL technologies do not use the whole bandwidth of the twisted pair, there is room remaining for a voice channel.

DSP digital signal processor. A DSP segments a voice signal into frames and stores them in voice packets.

DSS Digital Signature Standard. A digital signature algorithm designed by The US National Institute of Standards and Technology and based on public-key cryptography. DSS does not do user datagram encryption. DSS is a component in classic crypto, as well as the Redcreek IPsec card, but not in IPsec implemented in Cisco IOS software.

Dynamic NAT See NAT and address translation.

Dynamic PAT Dynamic Port Address Translation. Dynamic PAT lets multiple outbound sessions appear to originate from a single IP address. With PAT enabled, the ASA chooses a unique port number from the PAT IP address for each outbound translation slot (xlate). This feature is valuable when an ISP cannot allocate enough unique IP addresses for your outbound connections. The global pool addresses always come first, before a PAT address is used. See also NAT, Static PAT, and xlate.

E

ECHO See ping, ICMP. See also inspection engine.

EGP Exterior Gateway Protocol. Replaced by BGP. The ASA does not support EGP. See also BGP.

EIGRP Enhanced Interior Gateway Routing Protocol. The ASA does not support EIGRP.

EMBLEM Enterprise Management BaseLine Embedded Manageability. A syslog format designed to be consistent with the Cisco IOS system log format and is more compatible with CiscoWorks management applications.

encryption Application of a specific algorithm or cipher to data so as to render the data incomprehensible to those unauthorized to see the information. See also decryption.

ESMTP Extended SMTP. Extended version of SMTP that includes additional functionality, such as delivery notification and session delivery. ESMTP is described in RFC 1869, SMTP Service Extensions.

ESP Encapsulating Security Payload. An IPsec protocol, ESP provides authentication and encryption services for establishing a secure tunnel over an insecure network. For more information, refer to RFCs 2406 and 1827.

F

failover, failover mode Failover lets you configure two ASAs so that one will take over operation if the other one fails. The ASA supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover. With Active/Active failover, both units can pass network traffic. Active/Active failover lets you configure load balancing on your network. Active/Active failover is only available on units running in multiple context mode. With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.

Fixup See inspection engine.

Flash, Flash memory A nonvolatile storage device used to store the configuration file when the ASA is powered down.

FQDN/IP Fully qualified domain name/IP address. IPsec parameter that identifies peers that are security gateways.

FragGuard Provides IP fragment protection and performs full reassembly of all ICMP error messages and virtual reassembly of the remaining IP fragments that are routed through the ASA.

FTP File Transfer Protocol. Part of the TCP/IP protocol stack, used for transferring files between hosts.

G

GGSN gateway GPRS support node. A wireless gateway that allows mobile cell phone users to access the public data network or specified private IP networks.

global configuration mode Global configuration mode lets you change the ASA configuration. All user EXEC, privileged EXEC, and global configuration commands are available in this mode. See also user EXEC mode, privileged EXEC mode, command-specific configuration mode.

GMT Greenwich Mean Time. Replaced by UTC (Coordinated Universal Time) in 1967 as the world time standard.

GPRS general packet radio service. A service defined and standardized by the European Telecommunication Standards Institute. GPRS is an IP-packet-based extension of GSM networks and provides mobile, wireless, data communications

GRE Generic Routing Encapsulation described in RFCs 1701 and 1702. GRE is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to routers at remote points over an IP network. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single protocol backbone environment.

GSM Global System for Mobile Communication. A digital, mobile, radio standard developed for mobile, wireless, voice communications.

GTP GPRS tunneling protocol. GTP handles the flow of user packet data and signaling information between the SGSN and GGSN in a GPRS network. GTP is defined on both the Gn and Gp interfaces of a GPRS network.

H

H.225 A protocol used for TCP signaling in applications such as video conferencing. See also H.323 and inspection engine.

H.225.0 An ITU standard that governs H.225.0 session establishment and packetization. H.225.0 actually describes several different protocols: RAS, use of Q.931, and use of RTP.

H.245 An ITU standard that governs H.245 endpoint control.

H.320 Suite of ITU-T standard specifications for video conferencing over circuit-switched media, such as ISDN, fractional T-1, and switched-56 lines. Extensions of ITU-T standard H.320 enable video conferencing over LANs and other packet-switched networks, as well as video over the Internet.

H.323 Allows dissimilar communication devices to communicate with each other by using a standardized communication protocol. H.323 defines a common set of CODECs, call setup and negotiating procedures, and basic data transport methods.

H.323 RAS Registration, admission, and status signaling protocol. Enables devices to perform registration, admissions, bandwidth changes, and status and disengage procedures between VoIP gateway and the gatekeeper.

H.450.2 Call transfer supplementary service for H.323.

H.450.3 Call diversion supplementary service for H.323.

Hash, Hash Algorithm A hash algorithm is a one-way function that operates on a message of arbitrary length to create a fixed-length message digest used by cryptographic services to ensure its data integrity. MD5 has a smaller digest and is considered to be slightly faster than SHA-1. Cisco uses both SHA-1 and MD5 hashes within our implementation of the IPsec framework. See also encryption, HMAC, and VPN.

headend A firewall, concentrator, or other host that serves as the entry point into a private network for VPN client connections over the public network. See also ISP and VPN.

HMAC A mechanism for message authentication using cryptographic hashes such as SHA-1 and MD5.

host The name for any device on a TCP/IP network that has an IP address. See also network and node.

host/network An IP address and netmask used with other information to identify a single host or network subnet for ASA configuration, such as an address translation (xlate) or ACE.

HTTP Hypertext Transfer Protocol. A protocol used by browsers and web servers to transfer files. When a user views a web page, the browser can use HTTP to request and receive the files used by the web page. HTTP transmissions are not encrypted.

HTTPS Hypertext Transfer Protocol Secure. An SSL-encrypted version of HTTP.

I

IANA Internet Assigned Number Authority. Assigns all port and protocol numbers for use on the Internet.

ICMP Internet Control Message Protocol. Network-layer Internet protocol that reports errors and provides other information relevant to IP packet processing.

IDS Intrusion Detection System. A method of detecting malicious network activity by signatures and then implementing a policy for that signature.

IETF The Internet Engineering Task Force. A technical standards organization that develops RFC documents defining protocols for the Internet.

IGMP Internet Group Management Protocol. IGMP is a protocol used by IPv4 systems to report IP multicast memberships to neighboring multicast routers.

IKE Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPsec) that require keys. Before any IPsec traffic can be passed, each ASA must verify the identity of its peer. Identification can be done by manually entering preshared keys into both hosts or by a CA service. IKE is a hybrid protocol that uses part Oakley and part of another protocol suite called SKEME inside the ISAKMP framework. IKE (formerly known as ISAKMP/Oakley) is defined in RFC 2409.

IKE Extended Authentication IKE Extended Authenticate (Xauth) is implemented per the IETF draft-ietf-ipsec-isakmp-xauth-04.txt (extended authentication). This protocol provides the capability of authenticating a user within IKE using TACACS+ or RADIUS.

IKE Mode Configuration IKE Mode Configuration is implemented per the IETF draft-ietf-ipsec-isakmp-mode-cfg-04.txt. IKE Mode Configuration provides a method for a security gateway to download an IP address (and other network level configuration) to the VPN client as part of an IKE negotiation.

ILS Internet Locator Service. ILS is based on LDAP and is ILSv2 compliant. ILS was developed by Microsoft for use with its NetMeeting, SiteServer, and Active Directory products.

IMAP Internet Message Access Protocol. Method of accessing e-mail or bulletin board messages kept on a mail server that can be shared. IMAP permits client e-mail applications to access remote message stores as if they were local without actually transferring the message.

implicit rule An access rule automatically created by the ASA based on default rules or as a result of user-defined rules.

IMSI International Mobile Subscriber Identity. One of two components of a GTP tunnel ID, the other being the NSAPI. See also NSAPI.

inside The first interface, usually port 1, that connects your internal, trusted network protected by the ASA. See also interface, interface name.

inspection engine The ASA inspects certain application-level protocols to identify the location of embedded addressing information in traffic. Inspection allows NAT to translate these embedded addresses and to update any checksum or other fields that are affected by the translation. Because many protocols open secondary TCP or UDP ports, each application inspection engine also monitors sessions to determine the port numbers for secondary channels. 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. Some of the protocols that the ASA can inspect are CTIQBE, FTP, H.323, HTTP, MGCP, SMTP, and SNMP.

interface The physical connection between a particular network and a ASA.

interface IP address The IP address of the ASA network interface. Each interface IP address must be unique. Two or more interfaces must not be given the same IP address or IP addresses that are on the same IP network.

interface name Human-readable name assigned to the ASA network interface. The inside interface default name is “inside” and the outside interface default name is “outside.” See also inside and outside.

interface PAT The use of PAT where the PAT IP address is also the IP address of the outside interface. See Dynamic PAT, Static PAT.

Internet The global network that uses IP. Not a LAN. See also intranet.

intranet Intranetwork. A LAN that uses IP. See also network and Internet.

IP Internet Protocol. IP protocols are the most popular nonproprietary protocols because they can be used to communicate across any set of interconnected networks and are equally well suited for LAN and WAN communications.

IPS Intrusion Prevention Service. An in-line, deep-packet inspection-based solution that helps mitigate a wide range of network attacks.

IP address An IP protocol address. A ASA interface ip_address. IP version 4 addresses are 32 bits in length. This address space is used to designate the network number, optional subnetwork number, and a host number. The 32 bits are grouped into four octets (8 binary bits), represented by 4 decimal numbers separated by periods, or dots. The meaning of each of the four octets is determined by their use in a particular network.

IP pool A range of local IP addresses specified by a name, and a range with a starting IP address and an ending address. IP pools are used by DHCP and VPNs to assign local IP addresses to clients on the inside interface.

IPsec IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

IPsec Phase 1 The first phase of negotiating IPsec, includes the key exchange and the ISAKMP portions of IPsec.

IPsec Phase 2 The second phase of negotiating IPsec. Phase 2 determines the type of encryption rules used for payload, the source and destination that will be used for encryption, the definition of interesting traffic according to access lists, and the IPsec peer. IPsec is applied to the interface in Phase 2.

IPsec transform set A transform set specifies the IPsec protocol, encryption algorithm, and hash algorithm to use on traffic matching the IPsec policy. A transform describes a security protocol (AH or ESP) with its corresponding algorithms. The IPsec protocol used in almost all transform sets is ESP with the DES algorithm and HMAC-SHA for authentication.

ISAKMP Internet Security Association and Key Management Protocol. A protocol framework that defines payload formats, the mechanics of implementing a key exchange protocol, and the negotiation of a security association. See IKE.

ISP Internet Service Provider. An organization that provides connection to the Internet via their services, such as modem dial in over telephone voice lines or DSL.

J

JTAPI Java Telephony Application Programming Interface. A Java-based API supporting telephony functions. See also TAPI.

K

key A data object used for encryption, decryption, or authentication.

L

L2TP Layer Two Tunneling Protocol. An IETF standards track protocol defined in RFC 2661 that provides tunneling of PPP. L2TP is an extension to the PPP. L2TP merges the older Cisco Layer Two Forwarding (L2F) protocol with PPTP. L2TP can be used with IPsec encryption and is considered more secure against attack than PPTP.

LAN Local area network. A network residing in one location, such as a single building or campus. See also Internet, intranet, and network.

layer, layers Networking models implement layers with which different protocols are associated. The most common networking model is the OSI model, which consists of the following seven layers, in order: physical, data link, network, transport, session, presentation, and application.

LCN Logical channel number.

LDAP Lightweight Directory Access Protocol. LDAP provides management and browser applications with access to X.500 directories.

M

mask A 32-bit mask that shows how an Internet address is divided into network, subnet, and host parts. The mask has ones in the bit positions to be used for the network and subnet parts, and zeros for the host part. The mask should contain at least the standard network portion, and the subnet field should be contiguous with the network portion.

MCR See multicast.

MC router Multicast (MC) routers route multicast data transmissions to the hosts on each LAN in an internetwork that are registered to receive specific multimedia or other broadcasts. See also multicast.

MD5 Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and SHA-1 are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. SHA-1 is more secure than MD4 and MD5. Cisco uses hashes for authentication within the IPsec framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness. MD5 has a smaller digest and is considered to be slightly faster than SHA-1.

MDI media dependent interface.

MDIX media dependent interface crossover.

message digest A message digest is created by a hash algorithm, such as MD5 or SHA-1, that is used for ensuring message integrity.

MGCP Media Gateway Control Protocol. Media Gateway Control Protocol is a protocol for the control of VoIP calls by external call-control elements known as media gateway controllers or call agents. MGCP merges the IPDC and SGCP protocols.

Mode See Access Modes.

Mode Config See IKE Mode Configuration.

Modular Policy Framework A means of configuring ASA features in a manner similar to Cisco IOS software Modular QoS CLI.

MS mobile station. Refers generically to any mobile device, such as a mobile handset or computer, that is used to access network services. GPRS networks support three classes of MS, which describe the type of operation supported within the GPRS and the GSM mobile wireless networks. For example, a Class A MS supports simultaneous operation of GPRS and GSM services.

MS-CHAP Microsoft CHAP.

MTU maximum transmission unit. The maximum number of bytes in a packet that can flow efficiently across the network with best response time. For Ethernet, the default MTU is 1500 bytes, but each network can have different values, with serial connections having the smallest values. The MTU is described in RFC 1191.

multicast Refers to a network addressing method in which the source transmits a packet to multiple destinations, a multicast group, simultaneously. See also PIM, SMR.

N

N2H2 A third-party, policy-oriented filtering application that works with the ASA to control user web access. N2H2 can filter HTTP requests based on the destination hostname, destination IP address, username, and password. The N2H2 corporation was acquired by Secure Computing in October, 2003.

NAT Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into a globally routable address space.

NEM Network Extension Mode. Lets VPN hardware clients present a single, routable network to the remote private network over the VPN tunnel.

NetBIOS Network Basic Input/Output System. A Microsoft protocol that supports Windows hostname registration, session management, and data transfer. The ASA supports NetBIOS by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.

netmask See mask.

network In the context of ASA configuration, a network is a group of computing devices that share part of an IP address space and not a single host. A network consists of multiple nodes or hosts. See also host, Internet, intranet, IP, LAN, and node.

NMS network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer, such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources.

node Devices such as routers and printers that would not normally be called hosts. See also host, network.

nonvolatile storage, memory Storage or memory that, unlike RAM, retains its contents without power. Data in a nonvolatile storage device survives a power-off, power-on cycle.

NSAPI network service access point identifier. One of two components of a GTP tunnel ID, the other component being the IMSI. See also IMSI.

NSSA not-so-stubby-area. An OSPF feature described by RFC 1587. NSSA was first introduced in Cisco IOS software release 11.2. It is a nonproprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

NTLM NT Lan Manager. A Microsoft Windows challenge-response authentication method.

NTP Network Time Protocol.

O

Oakley A key exchange protocol that defines how to acquire authenticated keying material. The basic mechanism for Oakley is the Diffie-Hellman key exchange algorithm. Oakley is defined in RFC 2412.

object grouping Simplifies access control by letting you apply access control statements to groups of network objects, such as protocol, services, hosts, and networks.

OSPF Open Shortest Path First. OSPF is a routing protocol for IP networks. OSPF is a routing protocol widely deployed in large networks because of its efficient use of network bandwidth and its rapid convergence after changes in topology. The ASA supports OSPF.

OU Organizational Unit. An X.500 directory attribute.

outbound Refers to traffic whose destination is on an interface with lower security than the source interface.

outbound ACL An ACL applied to outbound traffic.

outside The first interface, usually port 0, that connects to other untrusted networks outside the ASA; the Internet. See also interface, interface name, outbound.

P

PAC PPTP Access Concentrator. A device attached to one or more PSTN or ISDN lines capable of PPP operation and of handling the PPTP protocol. The PAC needs to implement TCP/IP to pass traffic to one or more PNSs. It may also tunnel non-IP protocols.

PAT See Dynamic PAT, interface PAT, and Static PAT.

PDP Packet Data Protocol.

Perfmon The ASA feature that gathers and reports a wide variety of feature statistics, such as connections/second, xlates/second, and so on.

PFS Perfect Forwarding Secrecy. PFS enhances security by using a different security key for the IPsec Phase 1 and Phase 2 SAs. Without PFS, the same security key is used to establish SAs in both phases. PFS ensures that a given IPsec SA key 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 SA setup 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.

Phase 1 See IPsec Phase 1.

Phase 2 See IPsec Phase 2.

PIM Protocol Independent Multicast. PIM provides a scalable method for determining the best paths for distributing a specific multicast transmission to a group of hosts. Each host has registered using IGMP to receive the transmission. See also PIM-SM.

PIM-SM Protocol Independent Multicast-Sparse Mode. With PIM-SM, which is the default for Cisco routers, when the source of a multicast transmission begins broadcasting, the traffic is forwarded from one MC router to the next, until the packets reach every registered host. See also PIM.

ping An ICMP request sent by a host to determine if a second host is accessible.

PIX Private Internet eXchange. The Cisco PIX 500 series ASAs ranged from compact, plug-and-play desktop models for small/home offices to carrier-class gigabit models for the most demanding enterprise and service provider environments. Cisco PIX ASAs provided robust, enterprise-class integrated network security services to create a strong multilayered defense for fast changing network environments. The PIX has been replaced by the Cisco ASA 5500 series.

PKCS12 A standard for the transfer of PKI-related data, such as private keys, certificates, and other data. Devices supporting this standard let administrators maintain a single set of personal identity information.

PNS PPTP Network Server. A PNS is envisioned to operate on general-purpose computing/server platforms. The PNS handles the server side of PPTP. Because PPTP relies completely on TCP/IP and is independent of the interface hardware, the PNS may use any combination of IP interface hardware including LAN and WAN devices.

Policy NAT Lets you identify local traffic for address translation by specifying the source and destination addresses (or ports) in an access list.

POP Post Office Protocol. Protocol that client e-mail applications use to retrieve mail from a mail server.

Pool See IP pool.

Port A field in the packet headers of TCP and UDP protocols that identifies the higher level service which is the source or destination of the packet.

PPP Point-to-Point Protocol. Developed for dial-up ISP access using analog phone lines and modems.

PPPoE Point-to-Point Protocol over Ethernet. An IP protocol that encapsulates PPP packets and sends them over a local network or the internet to establish a connection to a host, usually between a client and an ISP.

PPTP Point-to-Point Tunneling Protocol. PPTP was introduced by Microsoft to provide secure remote access to Windows networks; however, because it is vulnerable to attack, PPTP is commonly used only when stronger security methods are not available or are not required. PPTP Ports are pptp, 1723/tcp, 1723/udp, and pptp. For more information about PPTP, see RFC 2637. See also PAC, PPTP GRE, PPTP GRE tunnel, PNS, PPTP session, and PPTP TCP.

PPTP GRE Version 1 of GRE for encapsulating PPP traffic.

PPTP GRE tunnel A tunnel defined by a PNS-PAC pair. The tunnel protocol is defined by a modified version of GRE. The tunnel carries PPP datagrams between the PAC and the PNS. Many sessions are multiplexed on a single tunnel. A control connection operating over TCP controls the establishment, release, and maintenance of sessions and of the tunnel itself.

PPTP session PPTP is connection-oriented. The PNS and PAC maintain the state for each user that is attached to a PAC. A session is created when an end-to-end PPP connection is attempted between a dial-up user and the PNS. The datagrams related to a session are sent over the tunnel between the PAC and PNS.

PPTP TCP Standard TCP session over which PPTP call control and management information is passed. The control session is logically associated with, but separate from, the sessions being tunneled through a PPTP tunnel.

preshared key A preshared key provides a method of IKE authentication that is suitable for networks with a limited, static number of IPsec peers. This method is limited in scalability because the key must be configured for each pair of IPsec peers. When a new IPsec peer is added to the network, the preshared key must be configured for every IPsec peer with which it communicates. Using certificates and CAs provides a more scalable method of IKE authentication.

primary, primary unit The ASA normally operating when two units, a primary and secondary, are operating in failover mode.

privileged EXEC mode The highest privilege level at the ASA CLI. Any user EXEC mode command will work in privileged EXEC mode. The privileged EXEC mode prompt appears as follows after you enter the **enable** command:

```
hostname> enable
hostname#
```

See also command-specific configuration mode, global configuration mode, user EXEC mode.

protocol, protocol literals A standard that defines the exchange of packets between network nodes for communication. Protocols work together in layers. Protocols are specified in the ASA configuration as part of defining a security policy by their literal values or port numbers. Possible ASA protocol literal values are ahp, eigrp, esp, gre, icmp, igmp, igrp, ip, ipinip, ipsec, nos, ospf, pcp, snp, tcp, and udp.

Proxy-ARP Enables the ASA to reply to an ARP request for IP addresses in the global pool. See also ARP.

public key A public key is one of a pair of keys that are generated by devices involved in public key infrastructure. Data encrypted with a public key can only be decrypted using the associated private key. When a private key is used to produce a digital signature, the receiver can use the public key of the sender to verify that the message was signed by the sender. These characteristics of key pairs provide a scalable and secure method of authentication over an insecure media, such as the Internet.

Q

QoS quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

R

RA Registration Authority. An authorized proxy for a CA. RAs can perform certificate enrollment and can issue CRLs. See also CA, certificate, public key.

RADIUS Remote Authentication Dial-In User Service. RADIUS is a distributed client/server system that secures networks against unauthorized access. RFC 2058 and RFC 2059 define the RADIUS protocol standard. See also AAA and TACACS+.

refresh Retrieve the running configuration from the ASA and update the screen. The icon and the button perform the same function.

registration authority See RA.

replay-detection A security service where the receiver can reject old or duplicate packets to defeat replay attacks. Replay attacks rely on the attacker sending out older or duplicate packets to the receiver and the receiver thinking that the bogus traffic is legitimate. Replay-detection is done by using sequence numbers combined with authentication and is a standard feature of IPsec.

RFC Request for Comments. RFC documents define protocols and standards for communications over the Internet. RFCs are developed and published by IETF.

RIP Routing Information Protocol. Interior Gateway Protocol (IGP) supplied with UNIX BSD systems. The most common IGP in the Internet. RIP uses hop count as a routing metric.

RLLA Reserved Link Local Address. Multicast addresses range from 224.0.0.0 to 239.255.255.255; however only the range 224.0.1.0 to 239.255.255.255 is available to users. The first part of the multicast address range, 224.0.0.0 to 224.0.0.255, is reserved and referred to as the RLLA. These addresses are unavailable.

route, routing The path through a network.

routed firewall mode In routed firewall mode, the ASA is counted as a router hop in the network. It performs NAT between connected networks and can use OSPF or RIP. See also transparent firewall mode.

RPC Remote Procedure Call. RPCs are procedure calls that are built or specified by clients and executed on servers, with the results returned over the network to the clients.

RSA A public key cryptographic algorithm (named after its inventors, Rivest, Shamir, and Adelman) with a variable key length. The main weakness of RSA is that it is significantly slow to compute compared to popular secret-key algorithms, such as DES. The Cisco implementation of IKE uses a Diffie-Hellman exchange to get the secret keys. This exchange can be authenticated with RSA (or preshared keys). With the Diffie-Hellman exchange, the DES key never crosses the network (not even in encrypted form), which is not the case with the RSA encrypt and sign technique. RSA is not public domain, and must be licensed from RSA Data Security.

RSH Remote Shell. A protocol that allows a user to execute commands on a remote system without having to log in to the system. For example, RSH can be used to remotely examine the status of a number of access servers without connecting to each communication server, executing the command, and then disconnecting from the communication server.

RTCP RTP Control Protocol. Protocol that monitors the QoS of an IPv6 RTP connection and conveys information about the ongoing session. See also RTP.

RTP Real-Time Transport Protocol. Commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, timestamping, and delivery monitoring to real-time applications.

RTSP Real Time Streaming Protocol. Enables the controlled delivery of real-time data, such as audio and video. RTSP is designed to work with established protocols, such as RTP and HTTP.

rule Conditional statements added to the ASA configuration to define security policy for a particular situation. See also ACE, ACL, NAT.

running configuration The configuration currently running in RAM on the ASA. The configuration that determines the operational characteristics of the ASA.

S

SA security association. An instance of security policy and keying material applied to a data flow. SAs are established in pairs by IPsec peers during both phases of IPsec. SAs specify the encryption algorithms and other security parameters used to create a secure tunnel. Phase 1 SAs (IKE SAs) establish a secure tunnel for negotiating Phase 2 SAs. Phase 2 SAs (IPsec SAs) establish the secure tunnel used for sending user data. Both IKE and IPsec use SAs, although SAs are independent of one another. IPsec SAs are unidirectional and they are unique in each security protocol. A set of SAs are needed for a protected data pipe, one per direction per protocol. For example, if you have a pipe that supports ESP between peers, one ESP SA is required for each direction. SAs are uniquely identified by destination (IPsec endpoint) address, security protocol (AH or ESP), and Security Parameter Index. IKE negotiates and establishes SAs on behalf of IPsec. A user can also establish IPsec SAs manually. An IKE SA is used by IKE only, and unlike the IPsec SA, it is bidirectional.

SCCP Skinny Client Control Protocol. A Cisco-proprietary protocol used between Cisco Call Manager and Cisco VoIP phones.

SCEP Simple Certificate Enrollment Protocol. A method of requesting and receiving (also known as enrolling) certificates from CAs.

SDP Session Definition Protocol. An IETF protocol for the definition of Multimedia Services. SDP messages can be part of SGCP and MGCP messages.

secondary unit The backup ASA when two are operating in failover mode.

secret key A secret key is a key shared only between the sender and receiver. See key, public key.

security context You can partition a single ASA into multiple virtual firewalls, known as security contexts. Each context is an independent firewall, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple stand-alone firewalls.

security services See cryptography.

serial transmission A method of data transmission in which the bits of a data character are transmitted sequentially over a single channel.

SGCP Simple Gateway Control Protocol. Controls VoIP gateways by an external call control element (called a call-agent).

SGSN Serving GPRS Support Node. The SGSN ensures mobility management, session management, and packet relaying functions.

SHA-1 Secure Hash Algorithm 1. SHA-1 [NIS94c] is a revision to SHA that was published in 1994. SHA is closely modeled after MD4 and produces a 160-bit digest. Because SHA produces a 160-bit digest, it is more resistant to brute-force attacks than 128-bit hashes (such as MD5), but it is slower. Secure Hash Algorithm 1 is a joint creation of the National Institute of Standards and Technology and the National Security Agency. This algorithm, like other hash algorithms, is used to generate a hash value, also known as a message digest, that acts like a CRC used in lower-layer protocols to ensure that message contents are not changed during transmission. SHA-1 is generally considered more secure than MD5.

SIP Session Initiation Protocol. Enables call handling sessions, particularly two-party audio conferences, or calls. SIP works with SDP for call signaling. SDP specifies the ports for the media stream. Using SIP, the ASA can support any SIP VoIP gateways and VoIP proxy servers.

site-to-site VPN A site-to-site VPN is established between two IPsec peers that connect remote networks into a single VPN. In this type of VPN, neither IPsec peer is the destination nor source of user traffic. Instead, each IPsec peer provides encryption and authentication services for hosts on the LANs connected to each IPsec peer. The hosts on each LAN send and receive data through the secure tunnel established by the pair of IPsec peers.

SKEME A key exchange protocol that defines how to derive authenticated keying material, with rapid key refreshment.

SMR Stub Multicast Routing. SMR allows the ASA to function as a stub router. A stub router is a device that acts as an IGMP proxy agent. IGMP is used to dynamically register specific hosts in a multicast group on a particular LAN with a multicast router. Multicast routers route multicast data transmissions to hosts that are registered to receive specific multimedia or other broadcasts. A stub router forwards IGMP messages between hosts and MC routers.

SMTP Simple Mail Transfer Protocol. SMTP is an Internet protocol that supports email services.

SNMP Simple Network Management Protocol. A standard method for managing network devices using data structures called Management Information Bases.

split tunneling Allows a remote VPN client simultaneous encrypted access to a private network and clear unencrypted access to the Internet. If you do not enable split tunneling, all traffic between the VPN client and the ASA is sent through an IPsec tunnel. All traffic originating from the VPN client is sent to the outside interface through a tunnel, and client access to the Internet from its remote site is denied.

spoofing A type of attack designed to foil network security mechanisms such as filters and access lists. A spoofing attack sends a packet that claims to be from an address from which it was not actually sent.

SQL*Net Structured Query Language Protocol. An Oracle protocol used to communicate between client and server processes.

SSC Security Services Card for the ASA 5505. For example, the AIP SSC.

SSH Secure Shell. An application running on top of a reliable transport layer, such as TCP/IP, that provides strong authentication and encryption capabilities.

SSL Secure Sockets Layer. A protocol that resides between the application layer and TCP/IP to provide transparent encryption of data traffic.

SSM Security Services Module. For example, the AIP SSM or CSC SSM.

standby unit See secondary unit.

stateful inspection Network protocols maintain certain data, called state information, at each end of a network connection between two hosts. State information is necessary to implement the features of a protocol, such as guaranteed packet delivery, data sequencing, flow control, and transaction or session IDs. Some of the protocol state information is sent in each packet while each protocol is being used. For example, a browser connected to a web server uses HTTP and supporting TCP/IP protocols. Each protocol layer maintains state information in the packets it sends and receives. The ASA and some other firewalls inspect the state information in each packet to verify that it is current and valid for every protocol it contains. This feature is called stateful inspection and is designed to create a powerful barrier to certain types of computer security threats.

Static PAT Static Port Address Translation. Static PAT is a static address that also maps a local port to a global port. See also Dynamic PAT, NAT.

subnetmask See mask.

T

TACACS+ Terminal Access Controller Access Control System Plus. A client-server protocol that supports AAA services, including command authorization. See also AAA, RADIUS.

TAPI Telephony Application Programming Interface. A programming interface in Microsoft Windows that supports telephony functions.

TCP Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission.

TCP Intercept With the TCP intercept feature, once the optional embryonic connection limit is reached, and until the embryonic connection count falls below this threshold, every SYN bound for the affected server is intercepted. For each SYN, the ASA responds on behalf of the server with an empty SYN/ACK segment. The ASA retains pertinent state information, drops the packet, and waits for the client acknowledgment. If the ACK is received, a copy of the client SYN segment is sent to the server and the TCP three-way handshake is performed between the ASA and the server. If this three-way handshake completes, the connection may resume as normal. If the client does not respond during any part of the connection phase, then the ASA retransmits the necessary segment using exponential back-offs.

TDP Tag Distribution Protocol. TDP is used by tag switching devices to distribute, request, and release tag binding information for multiple network layer protocols in a tag switching network. TDP does not replace routing protocols. Instead, it uses information learned from routing protocols to create tag bindings. TDP is also used to open, monitor, and close TDP sessions and to indicate errors that occur during those sessions. TDP operates over a connection-oriented transport layer protocol with guaranteed sequential delivery (such as TCP). The use of TDP does not preclude the use of other mechanisms to distribute tag binding information, such as piggybacking information on other protocols.

Telnet A terminal emulation protocol for TCP/IP networks such as the Internet. Telnet is a common way to control web servers remotely; however, its security vulnerabilities have led to its replacement by SSH.

TFTP Trivial File Transfer Protocol. TFTP is a simple protocol used to transfer files. It runs on UDP and is explained in depth in RFC 1350.

TID Tunnel Identifier.

TLS Transport Layer Security. A future IETF protocol to replace SSL.

traffic policing The traffic policing feature ensures that no traffic exceeds the maximum rate (bits per second) that you configure, which ensures that no one traffic flow can take over the entire resource.

transform set See IPsec transform set.

translate, translation See xlate.

transparent firewall mode A mode in which the ASA is not a router hop. You can use transparent firewall mode to simplify your network configuration or to make the ASA invisible to attackers. You can also use transparent firewall mode to allow traffic through that would otherwise be blocked in routed firewall mode. See also routed firewall mode.

transport mode An IPsec encryption mode that encrypts only the data portion (payload) of each packet but leaves the header untouched. Transport mode is less secure than tunnel mode.

TSP TAPI Service Provider. See also TAPI.

tunnel mode An IPsec encryption mode that encrypts both the header and data portion (payload) of each packet. Tunnel mode is more secure than transport mode.

tunnel A method of transporting data in one protocol by encapsulating it in another protocol. Tunneling is used for reasons of incompatibility, implementation simplification, or security. For example, a tunnel lets a remote VPN client have encrypted access to a private network.

Turbo ACL Increases ACL lookup speeds by compiling them into a set of lookup tables. Packet headers are used to access the tables in a small, fixed number of lookups, independent of the existing number of ACL entries.

U

UDP User Datagram Protocol. A connectionless transport layer protocol in the IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, which requires other protocols to handle error processing and retransmission. UDP is defined in RFC 768.

UMTS Universal Mobile Telecommunication System. An extension of GPRS networks that moves toward an all-IP network by delivering broadband information, including commerce and entertainment services, to mobile users via fixed, wireless, and satellite networks.

Unicast RPF Unicast Reverse Path Forwarding. Unicast RPF guards against spoofing by ensuring that packets have a source IP address that matches the correct source interface according to the routing table.

URL Uniform Resource Locator. A standardized addressing scheme for accessing hypertext documents and other services using a browser. For example, http://www.cisco.com.

user EXEC mode The lowest privilege level at the ASA CLI. The user EXEC mode prompt appears as follows when you first access the ASA:

```
hostname>
```

See also command-specific configuration mode, global configuration mode, and privileged EXEC mode.

UTC Coordinated Universal Time. The time zone at zero degrees longitude, previously called Greenwich Mean Time (GMT) and Zulu time. UTC replaced GMT in 1967 as the world time standard. UTC is based on an atomic time scale rather than an astronomical time scale.

UTRAN Universal Terrestrial Radio Access Network. Networking protocol used for implementing wireless networks in UMTS. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN, an SGSN and the UTRAN.

UUIE User-User Information Element. An element of an H.225 packet that identifies the users implicated in the message.

V

VLAN Virtual LAN. A group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same physical network cable, when they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.

VoIP Voice over IP. VoIP carries normal voice traffic, such as telephone calls and faxes, over an IP-based network. DSP segments the voice signal into frames, which are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323.

VPN Virtual Private Network. A network connection between two peers over the public network that is made private by strict authentication of users and the encryption of all data traffic. You can establish VPNs between clients, such as PCs, or a headend, such as the ASA.

virtual firewall See security context.

VSA Vendor-specific attribute. An attribute in a RADIUS packet that is defined by a vendor rather than by RADIUS RFCs. The RADIUS protocol uses IANA-assigned vendor numbers to help identify VSAs. This lets different vendors have VSAs of the same number. The combination of a vendor number and a VSA number makes a VSA unique. For example, the cisco-av-pair VSA is attribute 1 in the set of VSAs related to vendor number 9. Each vendor can define up to 256 VSAs. A RADIUS packet contains any VSAs attribute 26, named Vendor-specific. VSAs are sometimes referred to as subattributes.

W

WAN wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers.

WCCP Web Cache Communication Protocol. Transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times.

Websense A content filtering solution that manages employee access to the Internet. Websense uses a policy engine and a URL database to control user access to websites.

WEP Wired Equivalent Privacy. A security protocol for wireless LANs, defined in the IEEE 802.11b standard.

WINS Windows Internet Naming Service. A Windows system that determines the IP address associated with a particular network device, also known as name resolution. WINS uses a distributed database that is automatically updated with the NetBIOS names of network devices currently available and the IP address assigned to each one.WINS provides a distributed database for registering and querying dynamic NetBIOS names to IP address mapping in a routed network environment. It is the best choice for NetBIOS name resolution in such a routed network because it is designed to solve the problems that occur with name resolution in complex networks.

X

X.509 A widely used standard for defining digital certificates. X.509 is actually an ITU recommendation, which means that it has not yet been officially defined or approved for standardized usage.

xauth See IKE Extended Authentication.

xlate An xlate, also referred to as a translation entry, represents the mapping of one IP address to another, or the mapping of one IP address/port pair to another.

Symbols

/bits subnet masks A-3

Numerics

4GE SSM
connector types 15-12
fiber 15-12
SFP 15-12
802.1Q tagging 16-10
802.1Q trunk 15-34

A

AAA
about 46-1
accounting 53-17
addressing, configuring 88-5
authentication
CLI access 52-20
network access 53-2
proxy limit 53-11
authorization
command 52-23
downloadable access lists 53-13
network access 53-11
local database support 46-7
performance 53-1
server 96-4
adding 46-11, 46-12
types 46-1
support summary 46-3
web clients 53-8
AAA server group, add (group-policy) 87-6
ABR
definition of 33-2
access_rules 27-3
Access Control Server 86-31
Access Group pane
description 35-8
access lists
downloadable 53-13
global access rules 51-2
implicit deny 51-3
inbound 51-3
outbound 51-3
overview 51-1
phone proxy 65-7
access ports 16-8
access rules
turn off expansion 51-11
Accounting tab, tunnel group 87-91
ACE
add/edit/paste 87-32
Extended ACL tab 87-31
ACL
enabling IPSEC authenticated inbound sessions to bypass ACLs 87-103, 87-112
extended 87-31
for Clientless SSL VPN 87-40
standard 87-31
ACL Manager
Add/Edit/Paste ACE 87-32
dialog box 26-1, 87-31
activation key

entering 5-37
location 5-35
obtaining 5-36
Active/Active failover
about 13-1
actions 13-4
command replication 13-3
configuration synchronization 13-3
device initialization 13-3
duplicate MAC addresses, avoiding 13-2
optional settings
about 13-6
primary status 13-2
secondary status 13-2
triggers 13-4
Active/Standby failover
about 12-1
actions 12-4
command replication 12-3
configuration synchronization 12-2
device initialization 12-2
primary unit 12-2
secondary unit 12-2
triggers 12-4
Active Directory procedures B-15 to ??
Adaptive Security Algorithm 1-24
Add/Edit Access Group dialog box
description 35-8
Add/Edit Filtering Entry dialog box
description 33-19
Add/Edit IGMP Join Group dialog box
description 35-7
Add/Edit IGMP Static Group dialog box
description 35-8
Add/Edit Multicast Group dialog box 35-14
description 35-14
Add/Edit OSPF Area dialog box 33-13
description 33-13
Add/Edit OSPF Neighbor Entry dialog box 33-17, 33-28
description 33-17, 33-28
Add/Edit Rendezvous Point dialog box
restrictions 35-11
Add/Edit Summary Address dialog box
description 33-9, 33-13
Add/Edit Time Range dialog box 25-15
Add/Edit Virtual Link dialog box
description 33-20
add_acl 27-3
address assignment, client 87-92
Address Pool panel, VPN wizard 85-4
address pools, tunnel group 87-92
Address Translation Exemption panel, VPN wizard 85-6
admin context
about 9-2
changing 9-27
administrative access
using ICMP for 52-13
administrative distance 31-4
Advanced DHCP Options dialog box
description 23-5
Advanced OSPF Interface Properties dialog box 33-13
Advanced OSPF Virtual Link Properties dialog box
description 33-20
Advanced tab, tunnel group 87-92
ae_standard_access_list_rule 27-3
ae_webtype_acl 28-3
AIP
See IPS module
AIP SSC
loading an image 81-20, 81-22, 83-28
AIP SSM
about 81-1
loading an image 81-20, 81-22, 83-28
port-forwarding
enabling 17-7, 18-9
alternate address, ICMP message A-15
analyzing syslog messages 96-2
anti-replay window size 72-10, 86-12

APN, GTP application inspection 61-10
APPE command, denied request 58-19
application access
and e-mail proxy 90-7
and Web Access 90-7
configuring client applications 90-6
enabling cookies on browser 90-6
privileges 90-6
quitting properly 90-6
setting up on client 90-6
using e-mail 90-7
with IMAP client 90-7
Application Access Panel, WebVPN 91-78
application access using WebVPN
and hosts file errors 91-64
quitting properly 91-66
application firewall 58-27
application inspection
about 57-1
applying 57-6
configuring 57-6
security level requirements 17-2, 18-2, 19-2
Application Profile Customization Framework 91-15
Apply button 4-12
Area/Networks tab
description 33-6
area border router 33-2
ARP
NAT 39-22
ARP inspection
about 6-11
enabling 6-13
static entry 6-12
ARP spoofing 6-11
ARP table
monitoring 15-38, 16-12, 17-23, 18-22
ARP test, failover 11-19
ASA (Adaptive Security Algorithm) 1-24
ASA 5505
Base license 16-2
client
Xauth 87-106
MAC addresses 16-4
maximum VLANs 16-2
power over Ethernet 16-4
Security Plus license 16-2
SPAN 16-4
Spanning Tree Protocol, unsupported 16-8
ASA 5550 throughput 17-7, 18-9
ASA CX module
about 82-1
ASA feature compatibility 82-4
authentication proxy
about 82-3
port 82-12
troubleshooting 82-20
basic settings 82-9
cabling 82-7
configuration 82-6
failover 82-5
licensing 82-4
management access 82-2
management defaults 82-5
management IP address 82-8
monitoring 82-15
password reset 82-14
PRSM 82-3
reload 82-15
security policy 82-11
sending traffic to 82-12
shutdown 82-15
traffic flow 82-2
VPN 82-4
ASA CX Staus tab 4-30
ASBR
definition of 33-2
ASR groups 13-20
asymmetric routing

TCP state bypass 71-4
attacks
DNS HINFO request 78-9
DNS request for all records 78-10
DNS zone transfer 78-10
DNS zone transfer from high port 78-10
fragmented ICMP traffic 78-9
IP fragment 78-7
IP impossible packet 78-7
large ICMP traffic 78-9
ping of death 78-9
proxied RPC request 78-10
statd buffer overflow 78-11
TCP FIN only flags 78-9
TCP NULL flags 78-9
TCP SYN+FIN flags 78-9
UDP bomb 78-9
UDP chargen DoS 78-9
UDP snork 78-9
attributes
RADIUS B-26
Attributes Pushed to Client panel, VPN wizard 85-4
attribute-value pairs
TACACS+ B-37
authenticating a certificate 49-11, 49-13
authentication
about 46-2
CLI access 52-20
FTP 53-4
HTTP 53-3
network access 53-2
Telnet 53-3
web clients 53-8
WebVPN users with digital certificates 91-25
Authentication tab
description 33-11
Authentication tab, tunnel group 87-90
authorization
about 46-2
command 52-23
downloadable access lists 53-13
network access 53-11
Authorization tab, tunnel group 87-90
Auto-MDI/MDIX 15-2, 16-4
autostate messaging 2-10
Auto-Update, configuring 101-30

B

backed up configurations
restoring 101-25
backing up configurations 101-21
Backing Up the Local CA Server 101-24
Baltimore Technologies, CA server support 49-5
bandwidth 4-20
Basic tab
IPSec LAN-to-LAN, General tab 87-95
basic threat detection
See threat detection
bits subnet masks A-3
bookmarks
configuring to access Kerberos 91-45
Botnet Traffic Filter
actions 76-2
address categories 76-2
blacklist
adding entries 76-9
description 76-2
blocking traffic manually 76-12
classifying traffic 76-10
configuring 76-6
databases 76-2
default settings 76-6
DNS Reverse Lookup Cache
information about 76-4
maximum entries 76-4
using with dynamic database 76-9
DNS snooping 76-9

dropping traffic 76-11
graylist 76-11
dynamic database
enabling use of 76-7
files 76-3
information about 76-2
searching 76-13
updates 76-7
feature history 76-16
graylist
description 76-2
dropping traffic 76-11
guidelines and limitations 76-6
information about 76-1
licensing 76-6
monitoring 76-14
static database
adding entries 76-9
information about 76-3
syslog messages 76-14
task flow 76-7
threat level
dropping traffic 76-11
whitelist
adding entries 76-9
description 76-2
working overview 76-5
BPDUs
forwarding on the switch 2-10
broadcast Ping test 11-19
Browse ICMP 87-35
Browse Other 87-36
Browse Source or Destination Address 87-34
Browse Source or Destination Port 87-34
Browse Time Range 87-29
building blocks 25-1
bypassing firewall checks 71-3
bypassing the firewall, in the switch 2-7

C

CA
CRs and 49-3
public key cryptography 49-3
revoked certificates 49-3
supported servers 49-5
CA certificate 49-2
CA certificates 49-11, 49-13
call agents
MGCP application inspection 59-15, 59-16
Cancel button 4-12
CA server
Digicert 49-5
Geotrust 49-5
Godaddy 49-5
iPlanet 49-5
Netscape 49-5
RSA Keon 49-5
Thawte 49-5
Catalyst 6500
See switch
CDUP command, denied request 58-19
certificate
authentication, e-mail proxy 91-75
CA 49-11, 49-13
Cisco Unified Mobility 67-4
Cisco Unified Presence 68-4
code-signer 49-27
Identity 49-22
local CA 49-29
certificate authentication 49-11, 49-13
certificate enrollment 49-11, 49-13
Certificate Revocation Lists
See CRLs
change query interval 35-9
change query response time 35-9
change query timeout value 35-9
changing between contexts 9-25

changing the severity level 96-22
CIFS mount point
accessing 101-16
Cisco 7600
See switch
Cisco-AV-Pair LDAP attributes B-12
Cisco IOS CS CA
server support 49-5
Cisco IP Communicator 65-10
Cisco IP Phones, application inspection 59-32
Cisco UMA. See Cisco Unified Mobility.
Cisco Unified Mobility
architecture 67-2
ASA role 63-2, 63-3, 64-2
certificate 67-4
functionality 67-1
NAT and PAT requirements 67-3, 67-4
trust relationship 67-4
Cisco Unified Presence
ASA role 63-2, 63-3, 64-2
configuring the TLS Proxy 68-8
NAT and PAT requirements 68-2
trust relationship 68-4
Cisco UP. See Cisco Unified Presence.
Class A, B, and C addresses A-1
classes, logging
message class variables 96-4
types 96-4
classes, resource
See resource management
class map
regular expression 25-14
Client Address Assignment 87-92
Client Authentication panel, VPN wizard 85-3
Client Firewall tab 87-19
Clientless SSL VPN
client application requirements 90-2
client requirements 90-2
for file management 90-5
for network browsing 90-5
for web browsing 90-4
start-up 90-3
enable cookies for 90-6
end user set-up 90-1
printing and 90-3
remote requirements
for port forwarding 90-6
for using applications 90-6
remote system configuration and end-user requirements 90-3
security tips 90-2
supported applications 90-2
supported browsers 90-3
supported types of Internet connections 90-3
URL 90-3
username and password required 90-3
usernames and passwords 90-1
use suggestions 90-1
Client Update, edit , Windows and VPN 3002 clients 87-3
Client Update window, Windows and VPN 3002 clients 87-1
cluster
mixed scenarios 86-22
Cluster Dashboard tab 4-25
Cluster Firewall Dashboard tab 4-26
clustering
ASDM connection certificate IP address mismatch 10-10
backup owner 10-8
bootstrap configuration 10-44
cabling 10-25
cLACP
system ID 10-53
system priority 10-53
cluster control link
configuring 10-29
failure 10-7
overview 10-6
redundancy 10-7

size 10-6
configuration
examples 10-66
replication 10-8
configuring 10-51
connection
new, ownership 10-3
context mode 10-23
data path connection state replication 10-8
device-local EtherChannels, configuring on switch 10-22
executing a command cluster-wide 10-62
failover 10-23
feature history 10-71
features
centralized 10-15
individual units 10-16
NAT 10-19
SNMP 10-20
syslog and netflow 10-20
unsupported 10-15
VPN 10-20
guidelines and limitations 10-23
high availability 10-7
High Availability and Scalability Wizard 10-44
individual cluster interfaces, configuring 10-32
interface monitoring 10-7
IPv6 10-23
licensing 10-21
management
interface, configuring 10-32
interface, overview 10-9
network 10-9
overview 10-9
master unit
changing 10-61
election 10-3
maximum members 10-24
member requirements 10-24
model support 10-23
monitoring 10-63
overview
bootstrap configuration 10-3
cluster control link 10-6
Equal-Cost Multi-Path Routing 10-13
interfaces 10-4
load balancing 10-10
management 10-9
master unit 10-3
Policy-Based Routing 10-12
spanned EtherChannel 10-10
performance scaling factor 10-2
prerequisites 10-21
rebalancing new connections 10-14
removing a member 10-57
RSA key replication 10-10
software requirements 10-24
spanned EtherChannel
benefits 10-11
configuring 10-37
load balancing 10-11
maximum throughput 10-11
overview 10-10
redundancy 10-11
VSS or vPC 10-11
spanning-tree portfast 10-21
unit failure 10-8
unit health monitoring 10-7
upgrading software 10-24
code-signer certificate 49-27
command authorization
about 52-17
configuring 52-23
multiple contexts 52-18
configuration
factory default
commands 3-17
restoring 3-18

switch 2-1
configuration mode
accessing 3-2, 3-4
configurations, backing up 101-21
configuring
CSC activation 83-11
CSC email 83-21
CSC file transfer 83-22
CSC IP address 83-11
CSC license 83-11
CSC management access 83-13
CSC notifications 83-12
CSC password 83-13
CSC Setup Wizard 83-15, 83-18
CSC Setup Wizard Activation Codes Configuration 83-15
CSC Setup Wizard Host Configuration 83-16
CSC Setup Wizard IP Configuration 83-16
CSC Setup Wizard Management Access Configuration 83-17
CSC Setup Wizard Password Configuration 83-17
CSC Setup Wizard Summary 83-19
CSC Setup Wizard Traffic Selection for CSC Scan 83-17
CSC updates 83-23
CSC Web 83-20
configuring ASA
to join Active Directory domain 91-41
configuring bookmarks
to access Kerberos 91-45
configuring DNS 91-41
configuring mobile user security services 87-64
configuring MUS 87-64
connection limits
configuring 71-1
per context 9-18
console port logging 96-17
context mode 36-3
context modes 31-2, 32-3, 33-3, 34-3, 35-3, 83-6
contexts
See security contexts
conversion error, ICMP message A-16
creating a custom event list 96-17
CRL
cache refresh time 49-17, 49-21
CSC activation
configuring 83-11
CSC CPU
monitoring 83-27
CSC email
configuring 83-21
CSC file transfer
configuring 83-22
CSC IP address
configuring 83-11
CSC license
configuring 83-11
CSC management access
configuring 83-13
CSC memory
monitoring 83-27
CSC notifications
configuring 83-12
CSC password
configuring 83-13
CSC security events
monitoring 83-25
CSC Setup Wizard 83-15
activation codes configuratrion 83-15
Host configuratrion 83-16
IP configuratrion 83-16
management access configuratrion 83-17
password configuratrion 83-17
specifying traffic for CSC Scanning 83-18
summary 83-19
traffic selection for CSC Scan 83-17
CSC software updates
monitoring 83-26
CSC SSM

about 83-1
loading an image 81-20, 81-22, 83-28
what to scan 83-3
CSC SSM feature history 83-31
CSC SSM GUI
configuring 83-20
CSC threats
monitoring 83-24
CSC updates
configuring 83-23
CSC Web
configuring 83-20
customizing the end-user experience
by the security appliance 91-132
custom messages list
logging output destination 96-5
cut-through proxy
AAA performance 53-1
CX module
about 82-1
ASA feature compatibility 82-4
authentication proxy
about 82-3
port 82-12
troubleshooting 82-20
basic settings 82-9
cabling 82-7
configuration 82-6
failover 82-5
licensing 82-4
management access 82-2
management defaults 82-5
management IP address 82-8
monitoring 82-15
password reset 82-14
PRSM 82-3
reload 82-15
security policy 82-11
sending traffic to 82-12
shutdown 82-15
traffic flow 82-2
VPN 82-4

D

data flow
routed firewall 6-17
transparent firewall 6-23
date and time in messages 96-22
dead time 91-42
default
class 9-9
routes, defining equal cost routes 31-6
default configuration
commands 3-17
restoring 3-18
default policy 44-7
default routes
about 31-6
configuring 31-6
default tunnel gateway 87-4
destination address, browse 87-34
destination port, browse 87-34
device ID, including in messages 96-21
device ID in messages 96-21
Device Pass-Through 87-107
DHCP
addressing, configuring 88-6
configuring 23-3
monitoring
interface lease 17-24, 18-23
IP addresses 17-23, 18-23
server 17-23, 18-23
statistics 17-25, 18-24
statistics 17-25, 18-24
transparent firewall 51-5
DHCP relay
overview 23-6

DHCP Relay - Add/Edit DHCP Server dialog box
description 23-8
DHCP Relay pane
description 23-6
DHCP Relay panel 22-4, 23-6
DHCP Server pane
description 23-3
DHCP Server panel 23-3
DHCP services 21-4
DiffServ preservation 72-5
digital certificates 49-1, 49-2
authenticating WebVPN users 91-25
directory hierarchy search B-3
disabling content rewrite 91-31
disabling messages 96-22
disabling messages, specific message IDs 96-22
DMZ, definition 1-20
DNS
configuring 91-41
inspection
about 58-2
managing 58-1
rewrite, about 58-3
rewrite, configuring 58-3
NAT effect on 39-30
NAT effect on (8.2 and earlier) 42-13
server, configuring 21-8, 87-10
DNS HINFO request attack 78-9
DNS request for all records attack 78-10
DNS zone transfer attack 78-10
DNS zone transfer from high port attack 78-10
dotted decimal subnet masks A-3
downloadable access lists
configuring 53-13
converting netmask expressions 53-17
DSCP preservation 72-5
dual IP stack, configuring 17-2
dual-ISP support 31-8
duplex
interface 16-9, 16-12
duplex, configuring 15-12, 16-5
dynamic NAT
about 39-8
configuring (8.2 and earlier) 42-16
network object NAT 40-4
twice NAT 41-4
dynamic PAT
network object NAT 40-8
See also NAT
twice NAT 41-12

E

Easy VPN
client
Xauth 87-106
Easy VPN, advanced properties 87-107
Easy VPN client 87-105
Easy VPN Remote 87-105
echo reply, ICMP message A-15
ECMP 31-3
Edit DHCP Relay Agent Settings dialog box
description 23-7
Edit DHCP Server dialog box
description 23-4
Edit OSPF Interface Authentication dialog box 33-11
description 33-11
Edit OSPF Interface Properties dialog box 33-12
EIGRP 51-5
DUAL algorithm 36-2
hello interval 36-14
hello packets 36-1
hold time 36-2, 36-14
neighbor discovery 36-1
stub routing 36-5
stuck-in-active 36-2
e-mail
configuring for WebVPN 91-74

proxies, WebVPN 91-75
proxy, certificate authentication 91-75
WebVPN, configuring 91-74
e-mail proxy
and Clientless SSL VPN 90-7
enable command 3-2
Enable IPSec authenticated inbound sessions 87-103, 87-112
enabling logging 96-7
enabling secure logging 96-21
end-user interface, WebVPN, defining 91-77
enrolling
certificate 49-11, 49-13
Entrust, CA server support 49-5
established command, security level requirements 17-2, 18-2, 19-2
EtherChannel
adding interfaces 15-29
channel group 15-29
compatibility 15-5
converting existing interfaces 15-14
failover 15-10
guidelines 15-11
interface requirements 15-5
LACP 15-6
load balancing
configuring 15-31
overview 15-7
MAC address 15-8
management interface 15-28
maximum interfaces 15-31
minimum interfaces 15-31
mode
active 15-7
on 15-7
passive 15-7
overview 15-5
port priority 15-29
system priority 15-31
Ethernet
Auto-MDI/MDIX 15-2, 16-4
duplex 15-12, 16-5
jumbo frames, ASA 5580 15-37
jumbo frame support
single mode 17-15, 18-16
MTU 17-15, 18-16
speed 15-12, 16-5
EtherType access list
compatibilty with extended access lists 51-2
implicit deny 51-3
evaluation license 5-24
extended ACL 87-31
external group policy, configuring 87-6

F

factory default configuration
commands 3-17
restoring 3-18
failover
about 11-1
about virtual MAC addresses 12-15
Active/Active, *See* Active/Active failover
Active/Standby, *See* Active/Standby failover
configuration file
terminal messages, Active/Active 13-3
terminal messages, Active/Standby 12-2
contexts 12-2
criteria 12-14, 13-16
debug messages 11-21
defining standby IP addresses 12-12, 12-13
disabling 12-17, 13-21
enable 13-14
enabling Stateful Failover 12-11
Ethernet failover cable 11-3
failover link 11-3
forcing 12-16, 13-21
guidelines 83-6, 98-5

health monitoring 11-18
in multiple context mode 13-14
interface health 11-19
interface monitoring 11-19
interface tests 11-19
key 13-14
link communications 11-3
MAC addresses
about 12-2
automatically assigning 9-13
module placement
inter-chassis 11-12
intra-chassis 11-11
monitoring, health 11-18
network tests 11-19
primary unit 12-2
redundant interfaces 15-10
reset 13-23
restoring a failed group 12-17, 13-22
restoring a failed unit 12-17, 13-22
secondary unit 12-2
SNMP syslog traps 11-21
Stateful Failover 13-15
Stateful Failover, *See* Stateful Failover
state link 11-4
switch configuration 2-9
system log messages 11-20
system requirements 11-2
trunk 2-10
type selection 11-8
unit health 11-19
failover groups
about 13-16
adding 13-18
editing 13-18
monitoring 13-23
reset 13-23
fast path 1-25
fiber interfaces 15-12
Fibre Channel interfaces
default settings 27-2, 28-2, 51-7
filtering
rules 79-6
security level requirements 17-2, 18-2, 19-2
servers supported 79-2
URLs 79-1, 79-2
filtering messages 96-4
editing 96-27
Filtering pane
description 33-19
firewall, client, configuring settings 87-19
firewall mode
about 6-1
configuring 6-1
firewall server, Zone Labs 87-103
flash memory available for logs 96-20
flow control for 10 Gigabit Ethernet 15-23
flow-export actions 97-4
format of messages 96-3
fragmentation policy, IPsec 86-2
fragmented ICMP traffic attack 78-9
Fragment panel 78-2
fragment protection 1-22
fragment size 78-2
FTP
application inspection
viewing 58-8, 58-17, 58-28, 58-41, 58-48, 58-50, 59-7, 59-8, 59-15, 59-18, 59-26, 59-34, 59-35, 61-2, 61-12
filtering option 79-10
FTP inspection
about 58-12
configuring 58-12

G

gateway, default tunnel gateway 87-4
gateways
MGCP application inspection 59-16

global e-mail proxy attributes 91-75
graphs
bookmarking 15-40, 16-15, 17-28, 18-28
interface monitoring 15-40, 16-15, 17-28, 18-28
printing 15-40, 16-15, 17-28, 18-28
group policy
external, configuring 87-6
WINS and DNS servers 87-10
Group Policy window
add or edit, General tab 87-7, 87-23, 87-28
introduction 87-4
IPSec tab, add or edit 87-18
groups
SNMP 98-3
GTP
application inspection
viewing 61-6
GTP inspection
about 61-5
configuring 61-4

H

H.323
transparent firewall guidelines 6-4
H.323 inspection
about 59-3
configuring 59-2
limitations 59-4
Hardware Client tab 87-21
Help button 4-12
HELP command, denied request 58-19
Help menu 4-9
hierarchical policy, traffic shaping and priority queueing 72-11
high availability
about 11-1
history metrics 4-34
host
SNMP 98-3
hosts, subnet masks for A-3
hosts file
errors 91-64
reconfiguring 91-66
WebVPN 91-65
HSRP 6-4
HTTP
application inspection
viewing 58-27
filtering 79-1
configuring 79-9
HTTP(S)
filtering 79-2
HTTP inspection
about 58-21
configuring 58-21
HTTPS/Telnet/SSH
allowing network or host access to ASDM 52-1

I

ICMP
add group 87-36
browse 87-35
rules for access to ADSM 52-13
testing connectivity 73-1
type numbers A-15
ICMP Group 87-36
ICMP unreachable message limits 52-14
Identity Certificates 49-22
identity NAT
about 39-12
configuring (8.2 and earlier) 42-16
network object NAT 40-15
twice NAT 41-24
IKE Policy panel, VPN wizard 85-5
ILS inspection 60-1
IM 59-21

implementing SNMP 98-4
inbound access lists 51-3
individual syslog messages
  assigning or changing rate limits 96-23
information reply, ICMP message A-15
information request, ICMP message A-15
inside, definition 1-20
inspection engines
  See application inspection
installation
  module verification 2-3
Instant Messaging inspection 59-21
interface
  duplex 16-9, 16-12
  MTU 17-15, 18-16
  status 4-20
  subinterface, adding 15-36
  throughput 4-20
Interface pane 33-11
interfaces
  ASA 5505
    enabled status 16-8
    MAC addresses 16-4
    maximum VLANs 16-2
    switch port configuration 16-8
    trunk ports 16-10
  ASA 5550 throughput 17-7, 18-9
  default settings 27-2, 28-2, 51-7, 83-6
  duplex 15-12, 16-5
  failover monitoring 11-19
  fiber 15-12
  jumbo frame support
    single mode 17-15, 18-16
  MAC addresses
    automatically assigning 9-24
  monitoring 15-38, 16-13, 17-26, 18-26
  redundant 15-25
  SFP 15-12
  speed 15-12, 16-5
  subinterfaces 15-34
  turning off 17-22, 18-21
  turning on 17-22, 18-21
IOS
  upgrading 2-1
IP addresses
  classes A-1
  configuring an assignment method for remote access clients 88-1
  configuring for VPNs 88-1
  management, transparent firewall 18-8
  management, transparent firewall (8.3 and earlier) 19-4
  private A-2
  subnet mask A-4
IP audit
  enabling 78-5
  signatures 78-6
IP fragment attack 78-7
IP fragment database, displaying 78-2
IP fragment database, editing 78-3
IP impossible packet attack 78-7
IP overlapping fragments attack 78-7
IP phone
  phone proxy provisioning 65-11
IP phones
  addressing requirements for phone proxy 65-8
  supported for phone proxy 65-3, 66-3
IPS
  IP audit 78-5
IPSec
  anti-replay window 72-10
IPsec
  Cisco VPN Client 86-10
  fragmentation policy 86-2
IPSec rules
  anti-replay window size 72-10, 86-12
IPSec tab
  internal group policy 87-18

IPSec LAN-to-LAN 87-97
tunnel group 87-93
IPS module
about 81-1
configuration 81-7
operating modes 81-3
sending traffic to 81-18
traffic flow 81-2
virtual sensors 81-17
IP spoofing, preventing 78-1
IP teardrop attack 78-7
IPv6
configuring alongside IPv4 17-2
default route 31-7
dual IP stack 17-2
duplicate address detection 37-3
neighbor discovery 37-1
router advertisement messages 37-3
static neighbors 37-5
static routes 31-7
IPv6 addresses
anycast A-9
format A-5
multicast A-8
prefixes A-10
required A-10
types of A-6
unicast A-6
IPv6 prefixes 37-11
IPX 2-7

J

Java console 96-13
join Active Directory domain 91-41
Join Group pane
description 35-7
jumbo frames 17-14, 18-15
jumbo frames, ASA 5580 15-37
jumbo frame support
single mode 17-15, 18-16

K

KCD 91-38
Kerberos
configuring 46-11
support 46-6
Kerberos parameter 91-42
key pairs 49-23

L

LACP 15-6
large ICMP traffic attack 78-9
latency
about 72-1
configuring 72-2, 72-3
reducing 72-8
Layer 2 firewall
See transparent firewall
Layer 2 forwarding table
See MAC address table
Layer 3/4
matching multiple policy maps 44-5
LCS Federation Scenario 68-2
LDAP
application inspection 60-1
attribute mapping 46-19
Cisco-AV-pair B-12
configuring 46-11
configuring a AAA server B-2 to ??
directory search B-3
example configuration procedures B-15 to ??
hierarchy example B-3
SASL 46-6
user authentication 46-6

licenses
activation key
entering 5-37
location 5-35
obtaining 5-36
ASA 5505 5-2
ASA 5510 5-3, 5-8
ASA 5520 5-4
ASA 5540 5-5
ASA 5550 5-6
ASA 5580 5-7, 5-16
ASA 5585-X 5-13, 5-14, 5-15
Cisco Unified Communications Proxy features 63-4, 66-4, 67-6, 68-7, 69-8
default 5-24
evaluation 5-24
failover 5-34
guidelines 5-34
managing 5-1
preinstalled 5-24
Product Authorization Key 5-36
shared
backup server, information 5-28
client, configuring 5-39
communication issues 5-28
failover 5-29
maximum clients 5-30
monitoring 5-41
overview 5-27
server, configuring 5-38
SSL messages 5-28
temporary 5-24
viewing current 5-40
VPN Flex 5-24
licensing requirements
CSC SSM 83-5
logging 96-5
licensing requirements for SNMP 98-4
link up/down test 11-19
LLQ
See low-latency queue
load balancing
mixed cluster scenarios 86-22
local CA 49-29
Local CA User Database 49-33
local user database
adding a user 46-20
configuring 46-20
support 46-7
lockout recovery 52-30
logging
classes
filtering messages by 96-4
types 96-4
filtering
by message list 96-5
by severity level 96-1
output destinations
internal buffer 96-1, 96-7
Telnet or SSH session 96-7
queue
changing the size of 96-20
configuring 96-20
logging feature history 96-28
logging queue
configuring 96-20
login
banner, configuring 52-5
console 3-1
enable 3-2
FTP 53-4
global configuration mode 3-2
session 3-4
SSH 3-4, 52-4
Telnet 3-4
log viewers
executing certain commands 96-27
loops, avoiding 2-10

low-latency queue
applying 72-2, 72-3

M

MAC address
redundant interfaces 15-5
MAC addresses
ASA 5505 16-4
automatically assigning 9-24
failover 12-2
security context classification 9-3
MAC address table
about 6-23
built-in-switch 6-15
MAC learning, disabling 6-16
monitoring 15-38, 16-12, 17-26, 18-25
resource management 9-18
static entry 6-16
MAC learning, disabling 6-16
management interfaces
default settings 27-2, 28-2, 51-7
management IP address, transparent firewall 18-8
management IP address, transparent firewall (8.3 and earlier) 19-4
man-in-the-middle attack 6-11
mapped addresses
guidelines 39-21
guidelines (8.2 and earlier) 42-13
mask
reply, ICMP message A-15
request, ICMP message A-15
Master Passphrase 21-4
maximum sessions, IPSec 87-103
media termination address, criteria 65-6
menus 4-4
message filtering 96-4
message list
filtering by 96-5
messages, logging
classes
about 96-4
list of 96-4
component descriptions 96-3
filtering by message list 96-5
format of 96-3
severity levels 96-3
messages classes 96-4
messages in EMBLEM format 96-18, 96-19
metacharacters, regular expression 25-11
MGCP
application inspection
configuring 59-16
viewing 59-14
MGCP inspection
about 59-12
configuring 59-12
mgmt0 interfaces
default settings 27-2, 28-2, 51-7
MIBs for SNMP 98-11
Microsoft Access Proxy 68-1
Microsoft KCD 91-38
Microsoft Windows CA, supported 49-5
mixed cluster scenarios, load balancing 86-22
MMP inspection 67-1
mobile redirection, ICMP message A-16
mode
context 9-16
firewall 6-1
monitoring
ARP table 15-38, 16-12, 17-23, 18-22
CSC CPU 83-27
CSC memory 83-27
CSC security events 83-25
CSC software updates 83-26
CSC SSM 83-24
CSC threats 83-24
DHCP

interface lease 17-24, 18-23
IP addresses 17-23, 18-23
server 17-23, 18-23
statistics 17-25, 18-24
failover 11-18
failover groups 13-23
history metrics 4-34
interfaces 15-38, 16-13, 17-26, 18-26
MAC address table 15-38, 16-12, 17-26, 18-25
OSPF 33-35
SNMP 98-1
monitoring logging 96-25
monitoring NSEL 97-7
monitoring switch traffic, ASA 5505 16-4
MPF
default policy 44-7
feature directionality 44-3
features 44-1
flows 44-5
matching multiple policy maps 44-5
See also class map
See also policy map
MPLS
LDP 51-6
router-id 51-6
TDP 51-6
MRoute pane
description 35-5
MSFC
overview 1-19
SVIs 2-7
MTU 17-15, 18-16
multicast traffic 6-4
multiple context mode
logging 96-2
See security contexts
multiple SVIs 2-6
multi-session PAT 40-18
MUS
configuring 87-64

N

NAT
about 39-1, 42-1
about (8.2 and earlier) 42-1
bidirectional initiation 39-2
bypassing NAT (8.2 and earlier) 42-10
disabling proxy ARP for global addresses 30-11
DNS 39-30
DNS (8.2 and earlier) 42-13
dynamic
about 39-8
dynamic NAT
about (8.2 and earlier) 42-6
configuring (8.2 and earlier) 42-22
implementation (8.2 and earlier) 42-16
network object NAT 40-4
twice NAT 41-4
dynamic PAT
about 39-10
network object NAT 40-8
twice NAT 41-12
exemption (8.2 and earlier) 42-10
identity
about 39-12
identity NAT
about (8.2 and earlier) 42-10
network object NAT 40-15
twice NAT 41-24
implementation 39-15
interfaces 39-21
mapped address guidelines 39-21
network object
comparison with twice NAT 39-15
network object NAT
about 39-16
configuring 40-1

dynamic NAT 40-4
dynamic PAT 40-8
examples 40-20
guidelines 40-2
identity NAT 40-15
monitoring 40-19
prerequisites 40-2
static NAT 40-11
no proxy ARP 40-17
object
extended PAT 40-4
flat range for PAT 40-4
PAT
about (8.2 and earlier) 42-8
configuring (8.2 and earlier) 42-22
implementation (8.2 and earlier) 42-16
policy NAT, about (8.2 and earlier) 42-10
routed mode 39-13
route lookup 40-17, 41-29
RPC not supported with 60-3
rule order 39-20
rule order (8.2 and earlier) 42-13
same security level (8.2 and earlier) 42-12
static
about 39-3
few-to-many mapping 39-7
many-to-few mapping 39-6, 39-7
one-to-many 39-6
static NAT
about (8.2 and earlier) 42-8
configuring (8.2 and earlier) 42-26
network object NAT 40-11
twice NAT 41-18
static PAT
about (8.2 and earlier) 42-9
static with port translation
about 39-4
terminology 39-2
transparent mode 39-13
transparent mode (8.2 and earlier) 42-3
twice
extended PAT 41-4
flat range for PAT 41-4
twice NAT
about 39-16
comparison with network object NAT 39-15
configuring 41-1
dynamic NAT 41-4
dynamic PAT 41-12
examples 41-30
guidelines 41-2
identity NAT 41-24
monitoring 41-30
prerequisites 41-2
static NAT 41-18
types 39-3
types (8.2 and earlier) 42-6
VPN 39-24
VPN client rules 39-20
neighbor reachable time 37-3
neighbor solicitation messages 37-2
neighrbor advertisement messages 37-2
NetBIOS server
tab 87-72
NetFlow
overview 97-1
NetFlow event
matching to configured collectors 97-6
Network Activity test 11-19
Network Admission Control
uses, requirements, and limitations 86-30
network object NAT
about 39-16
comparison with twice NAT 39-15
configuring 40-1
dynamic NAT 40-4
dynamic PAT 40-8
examples 40-20

guidelines 40-2
identity NAT 40-15
monitoring 40-19
prerequisites 40-2
static NAT 40-11
No Payload Encryption 5-33
NSEL and syslog messages
redundant messages 97-2
NSEL feature history 97-8
NSEL licensing requirements 97-4
NTLM support 46-6
NT server
configuring 46-11
support 46-6

O

object NAT
See network object NAT
open ports A-14
Options menu 4-5
OSPF
area parameters 33-13
authentication support 33-2
configuring authentication 33-11
defining a static neighbor 33-17, 33-28
defining interface properties 33-12
interaction with NAT 33-2
interface parameters 33-11
interface properties 33-11, 33-12
link-state advertisement 33-2
logging neighbor states 33-18
LSAs 33-2
monitoring 33-35
NSSA 33-14
processes 33-2
redistributing routes 33-6
route calculation timers 33-18
route summarization 33-10
OSPF parameters
dead interval 33-13
hello interval 33-13
retransmit interval 33-13
transmit delay 33-13
outbound access lists 51-3
Outlook Web Access (OWA) and Clientless SSL VPN 90-7
output destination 96-6
output destinations 96-1, 96-7
e-mail address 96-1, 96-7
SNMP management station 96-1, 96-7
Telnet or SSH session 96-1, 96-7
outside, definition 1-20
oversubscribing resources 9-10

P

packet
classifier 9-3
packet flow
routed firewall 6-17
transparent firewall 6-23
packet trace, enabling 73-7
parameter problem, ICMP message A-15
password
Clientless SSL VPN 90-1
passwords
WebVPN 91-123
PAT
per-session and multi-session 40-18
See dynamic PAT
PAT pool 40-7, 41-9
round robin 40-7, 41-9
pause frames for flow control 15-23
PDA support for WebVPN 91-74
PDP context, GTP application inspection 61-8
per-session PAT 40-18
phone proxy

access lists 65-7
ASA role 63-3
Cisco IP Communicator 65-10
Cisco UCM supported versions 65-3, 66-3
IP phone addressing 65-8
IP phone provisioning 65-11
IP phones supported 65-3, 66-3
Linksys routers, configuring 65-21
NAT and PAT requirements 65-8
ports 65-7
rate limiting 65-10
TLS Proxy on ASA, described 63-3
PIM
shortest path tree settings 35-13
ping
See ICMP
using 73-3
ping of death attack 78-9
PoE 16-4
policy, QoS 72-1
policy map
Layer 3/4
about 44-1
feature directionality 44-3
flows 44-5
policy NAT, about (8.2 and earlier) 42-10
Port Forwarding
configuring client applications 90-6
port-forwarding
enabling 17-7, 18-9
ports
open on device A-14
phone proxy 65-7
TCP and UDP A-11
port translation
about 39-4
posture validation
uses, requirements, and limitations 86-30
Posture Validation Exception, add/edit 86-32
power over Ethernet 16-4
PPP tab, tunnel-group 87-95
prerequisites for use
CSC SSM 83-5
presence_proxy_remotecert 64-15
primary unit, failover 12-2
printing
graphs 15-40, 16-15, 17-28, 18-28
priority queueing
hierarchical policy with traffic shaping 72-11
IPSec anti-replay window size 72-10, 86-12
private networks A-2
privileged EXEC mode
accessing 3-4
privileged EXEC mode, accessing 3-2
privileged mode
accessing 3-2
Process Instances tab
description 33-5
Product Authorization Key 5-36
Properties tab 33-12
description 33-12
fields 33-12
Protocol Group, add 87-37
protocol numbers and literal values A-11
Protocol pane (PIM)
description 35-10
proxied RPC request attack 78-10
proxy
See e-mail proxy
proxy ARP
NAT
NAT
proxy ARP 39-22
proxy ARP, disabling 30-11
proxy bypass 91-76
proxy servers
SIP and 59-20
PRSM 82-3

public key cryptography 49-3

Q

QoS
about 72-1, 72-3
DiffServ preservation 72-5
DSCP preservation 72-5
feature interaction 72-4
policies 72-1
priority queueing
hierarchical policy with traffic shaping 72-11
IPSec anti-replay window 72-10
IPSec anti-replay window size 72-10, 86-12
statistics 72-11
token bucket 72-2
traffic shaping
overview 72-4
viewing statistics 72-11, 72-12
Quality of Service
See QoS
queue, logging
changing the size of 96-20
queue, QoS
latency, reducing 72-8
limit 72-2, 72-3

R

RADIUS
attributes B-26
Cisco AV pair B-12
configuring a AAA server B-25
configuring a server 46-11
downloadable access lists 53-13
network access authentication 53-6
network access authorization 53-13
support 46-4
rapid link failure detection 2-10
rate limit 96-22
rate limiting 72-3
rate limiting, phone proxy 65-10
reactivation mode 91-42
realm 91-42
RealPlayer 59-17
recurring time range, add or edit 87-30
redirect, ICMP message A-15
redundant interface
EtherChannel
converting existing interfaces 15-14
redundant interfaces
configuring 15-25
failover 15-10
MAC address 15-5
setting the active interface 15-28
Registration Authority description 49-3
regular expression 25-10
reloading
context 9-29
Request Filter pane
description 35-12
reset
inbound connections 78-3
outside connections 78-3
Reset button 4-12
resetting the services module 2-11
resource management
about 9-10
class 9-17
configuring 9-8
default class 9-9
oversubscribing 9-10
resource types 9-18
unlimited 9-11
restoring backups 101-25
revoked certificates 49-3
rewrite, disabling 91-31

RFCs for SNMP 98-11
RIP
authentication 34-2
definition of 34-1
enabling 34-4
support for 34-2
RIP panel
limitations 34-3
RIP Version 2 Notes 34-3
RNFR command, denied request 58-20
RNTO command, denied request 58-20
routed mode
about 6-1
NAT 39-13
setting 6-1
route map
definition 32-1
route maps
defining 32-4
uses 32-1
router
advertisement, ICMP message A-15
solicitation, ICMP message A-15
router advertisement messages 37-3
router advertisement transmission interval 37-9
router lifetime value 37-9
routes
about default 31-6
configuring default routes 31-6
configuring IPv6 default 31-7
configuring IPv6 static 31-7
configuring static routes 31-3
Route Summarization tab
description 33-6
Route Tree pane 35-13
description 35-13
routing
other protocols 51-5
RTSP inspection
about 59-17
configuring 59-16
rules
ICMP 52-12
running configuration
copying 101-20

S

same security level communication
enabling 17-20, 18-21, 19-17
NAT (8.2 and earlier) 42-12
SCCP (Skinny) inspection
about 59-32
configuration 59-32
configuring 59-31
SDI
configuring 46-11
support 46-5
secondary unit, failover 12-2
Secure Computing SmartFilter filtering server 79-3
Secure Copy
configure server 52-9
security appliance
connecting to 3-1
managing licenses 5-1
security contexts
about 9-1
adding 9-20
admin context
about 9-2
changing 9-27
cascading 9-6
changing between 9-25
classifier 9-3
command authorization 52-18
configuration
URL, changing 9-28
logging in 9-7

MAC addresses
automatically assigning 9-24
classifying using 9-3
managing 9-1, 9-26
monitoring 9-30
MSFC compatibility 1-20
multiple mode, enabling 9-16
nesting or cascading 9-7
reloading 9-29
removing 9-26
resource management 9-10
unsupported features 9-15
security level
about 17-1
security models for SNMP 98-3
segment size
maximum and minimum 78-3
sending messages to an e-mail address 96-13
sending messages to a specified output destination 96-21
sending messages to a syslog server 96-8
sending messages to a Telnet or SSH session 96-17
sending messages to the console port 96-17
sending messages to the internal log buffer 96-12
Server and URL List
add/edit 87-41
server group name 91-42
Server or URL
dialog box 87-42
session management path 1-24
severity levels, of system log messages
changing 96-1
filtering by 96-1
list of 96-3
severity levels, of system messages
definition 96-3
shared license
backup server, information 5-28
client, configuring 5-39
communication issues 5-28
failover 5-29
maximum clients 5-30
monitoring 5-41
server, configuring 5-38
SSL messages 5-28
shun
duration 77-10
signatures
attack and informational 78-6
single mode
backing up configuration 9-16
configuration 9-16
enabling 9-16
restoring 9-17
SIP inspection
about 59-20
configuring 59-20
instant messaging 59-21
SITE command, denied request 58-20
Smart Call Home monitoring 99-10
smart tunnels 91-45
SMTP inspection 58-47
SNMP
about 98-1
application inspection
viewing 61-14
failover 98-5
management station 96-1, 96-7
prerequisites 98-4
SNMP configuration 98-6
SNMP groups 98-3
SNMP hosts 98-3
SNMP management station
adding 98-6
SNMP monitoring 98-9, 98-10
SNMP terminology 98-2
SNMP users 98-3
SNMP Version 3 98-3, 98-8
SNMP Versions 1 and 2c 98-8

software
version 4-30
source address, browse 87-34
source port, browse 87-34
source quench, ICMP message A-15
SPAN 16-4
Spanning Tree Protocol, unsupported 16-8
SPAN session 2-2
specifying traffic for CSC scanning 83-18
speed, configuring 15-12, 16-5
SSCs
management access 81-4
management defaults 81-6
management interface 81-14
password reset 81-22, 83-29
reload 81-23, 83-30
reset 81-23, 83-30
routing 81-10
sessioning to 81-13
shutdown 81-22, 83-30
SSH
concurrent connections 52-2
login 52-4
username 52-4
SSMs
loading an image 81-20, 81-22, 83-28
management access 81-4
management defaults 81-6
password reset 81-22, 83-29
reload 81-23, 83-30
reset 81-23, 83-30
routing 81-10
sessioning to 81-13
shutdown 81-22, 83-30
SSO with WebVPN
configuring HTTP form protocol 91-22
Standard Access List Rule, add/edit 87-41
Standard ACL tab 87-31
startup configuration
copying 101-20
Startup Wizard
acessing 7-1
licensing requirements 7-1, 64-3
statd buffer overflow attack 78-11
Stateful Failover
about 11-10
enabling 12-11
settings 13-15
state information 11-10
state link 11-4
stateful inspection 1-24
bypassing 71-3
state information 11-10
state link 11-4
static ARP entry 6-12
static bridge entry 6-16
Static Group pane
description 35-7
static NAT
about 39-3
few-to-many mapping 39-7
many-to-few mapping 39-6, 39-7
network object NAT 40-11
twice NAT 41-18
static NAT with port translation
about 39-4
static PAT
See PAT
static routes
configuring 31-3
deleting 31-6
statistics, QoS 72-11
status bar 4-11
stealth firewall
See transparent firewall
STOU command, denied request 58-20
stuck-in-active 36-2
subinterface

adding 15-36
subinterfaces, adding 15-34
subnet masks
/bits A-3
about A-2
address range A-4
determining A-3
dotted decimal A-3
number of hosts A-3
subordinate certificate 49-2
Summary Address pane
description 33-9
Summary panel, VPN wizard 85-6
Sun Microsystems Java™ Runtime Environment (JRE) and Clientless SSL VPN 90-6
Sun RPC inspection
about 60-3
configuring 60-3
SVIs
configuring 2-8
multiple 2-6
overview 2-6
switch
assigning VLANs to module 2-4
autostate messaging 2-10
BPDU forwarding 2-10
configuration 2-1
failover compatibility with transparent firewall 2-10
failover configuration 2-9
trunk for failover 2-10
verifying module installation 2-3
switched virtual interfaces
See SVIs
switch MAC address table 6-15
switch ports
access ports 16-8
SPAN 16-4
trunk ports 16-10
syslogd server program 96-6
syslog message filtering
using log viewers 96-25
syslog messages
analyzing 96-2
syslog messaging for SNMP 98-10
syslog server
designating more than one as output destination 96-6
EMBLEM format
configuring 96-19
system configuration 9-2
system log messages
classes 96-4
classes of 96-4
configuring in groups
by message list 96-5
by severity level 96-1
device ID, including 96-21
disabling logging of 96-1
filtering by message class 96-4
output destinations 96-1, 96-7
syslog message server 96-7
Telnet or SSH session 96-7
severity levels
about 96-3
changing the severity level of a message 96-1
timestamp, including 96-22

T

TACACS+
command authorization, configuring 52-28
configuring a server 46-11
network access authorization 53-11
support 46-5
tail drop 72-3
TCP
connection limits per context 9-18
maximum segment size 78-3
ports and literal values A-11

TIME_WAIT state 78-4
TCP FIN only flags attack 78-9
TCP Intercept
statistics 77-6
TCP normalization 71-3
TCP NULL flags attack 78-9
TCP Service Group, add 87-35
TCP state bypass
AAA 71-5
configuring 71-8
failover 71-5
firewall mode 71-5
inspection 71-5
mutliple context mode 71-5
NAT 71-5
SSMs and SSCs 71-5
TCP Intercept 71-5
TCP normalization 71-5
unsupported features 71-5
TCP SYN+FIN flags attack 78-9
Telnet
allowing management access 52-1
concurrent connections 52-2
login 52-4
temporary license 5-24
testing configuration 73-1
threat detection
basic
drop types 77-2
enabling 77-4
overview 77-2
rate intervals 77-2
statistics, viewing 77-4
system performance 77-2
scanning
enabling 77-10
host database 77-9
overview 77-8
shunning attackers 77-10
system performance 77-9
scanning statistics
enabling 77-6
system performance 4-23, 77-5
viewing 77-7
shun
duration 77-10
TIME_WAIT state 78-4
time exceeded, ICMP message A-15
time range
add or edit 87-29
browse 87-29
recurring 87-30
timestamp, including in system log messages 96-22
timestamp reply, ICMP message A-15
timestamp request, ICMP message A-15
TLS Proxy
applications supported by ASA 63-3
Cisco Unified Presence architecture 68-1
configuring for Cisco Unified Presence 68-8
licenses 63-4, 66-4, 67-6, 68-7, 69-8
tocken bucket 72-2
toolbar, floating, WebVPN 91-79
Tools menu 4-6
traceroute, enabling 4-7, 73-6
traffic flow
routed firewall 6-17
transparent firewall 6-23
traffic shaping
overview 72-4
transmit queue ring limit 72-2, 72-3
transparent firewall
about 6-2
ARP inspection
about 6-11
enabling 6-13
static entry 6-12
data flow 6-23
DHCP packets, allowing 51-5

guidelines 6-7
H.323 guidelines 6-4
HSRP 6-4
MAC learning, disabling 6-16
management IP address 18-8
management IP address (8.3 and earlier) 19-4
multicast traffic 6-4
packet handling 51-5
static bridge entry 6-16
unsupported features 6-8
VRRP 6-4
transparent mode
NAT 39-13
NAT (8.2 and earlier) 42-3
trunk, 802.1Q 15-34
trunk ports 16-10
Trusted Flow Acceleration
modes 6-6, 6-11, 6-15, 7-1, 13-7, 26-2, 51-7
trustpoint 49-4
trust relationship
Cisco Unified Mobility 67-4
Cisco Unified Presence 68-4
Tunneled Management 87-107
tunnel gateway, default 87-4
twice NAT
about 39-16
comparison with network object NAT 39-15
configuring 41-1
dynamic NAT 41-4
dynamic PAT 41-12
examples 41-30
guidelines 41-2
identity NAT 41-24
monitoring 41-30
prerequisites 41-2
static NAT 41-18
tx-ring-limit 72-2, 72-3

U

UDP
bomb attack 78-9
chargen DoS attack 78-9
connection limits per context 9-18
connection state information 1-24
ports and literal values A-11
snork attack 78-9
unprivileged mode
accessing 3-4
unreachable, ICMP message A-15
unreachable messages
required for MTU discovery 52-12
upgrading
IOS 2-1
URL
filtering
configuring 79-9
URLs
context configuration, changing 9-28
filtering 79-1
filtering, about 79-2
User Accounts panel, VPN wizard 85-4
user EXEC mode
accessing 3-1
username
adding 46-20
Clientless SSL VPN 90-1
WebVPN 91-123
Xauth for Easy VPN client 87-106
users
SNMP 98-3
using clustering 96-5, 97-3

V

VeriSign, configuring CAs example 49-6
version

IPS software 4-30
viewing QoS statistics 72-11, 72-12
virtual firewalls
See security contexts
virtual HTTP 53-3
Virtual Link
description 33-20
virtual MAC address
defining for Active/Active failover 13-19
virtual MAC addresses
about 12-15, 13-19
defaults for Active/Active failover 13-19
defining for Active/Standby failover 13-20
virtual private network
overview 85-1
virtual reassembly 1-22
virtual sensors 81-17
VLANs 15-34
802.1Q trunk 15-34
ASA 5505
MAC addresses 16-4
maximum 16-2
assigning to FWSM 2-4
interfaces 2-4
subinterfaces 15-34
VoIP
proxy servers 59-20
VPN
address range, subnets A-4
overview 85-1
system options 87-102
VPN client
NAT rules 39-20
VPN Client, IPsec attributes 86-10
VPN flex license 5-24
VPN Tunnel Type panel, VPN wizard 85-2
VPN wizard 85-1
Address Pool panel 85-4
Address Translation Exemption panel 85-6
Attributes Pushed to Client panel 85-4
Client Authentication panel 85-3
IKE Policy panel 85-5
Summary panel 85-6
User Accounts panel 85-4
VPN Tunnel Type panel 85-2
VRRP 6-4

W

WCCP 54-1
web browsing with Clientless SSL VPN 90-4
web caching 54-1
web clients, secure authentication 53-8
web e-Mail (Outlook Web Access), Outlook Web Access 91-75
Websense filtering server 79-3
WebVPN
authenticating with digital certificates 91-25
client application requirements 91-123
client requirements 91-123
configuring
e-mail 91-74
defining the end-user interface 91-77
definition 91-2
e-mail 91-74
e-mail proxies 91-75
end user set-up 91-77
floating toolbar 91-79
hosts file 91-65
hosts files, reconfiguring 91-66
PDA support 91-74
security tips 91-123
supported applications 91-123
troubleshooting 91-64
usernames and passwords 91-123
use suggestions 90-2, 91-77, 91-123
WebVPN, Application Access Panel 91-78
Window menu 4-9

Windows Service Account
adding in Active Directory 91-40
WINS server, configuring 87-10
Wizards menu 4-8

X

Xauth, Easy VPN client 87-106
XOFF frames 15-23

Z

Zone Labs Integrity Server 87-103