

Cisco ASA Series General Operations ASDM Configuration Guide

Software Version 7.1
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: December 3, 2012
Updated: March 31, 2014

Cisco Systems, Inc.
www.cisco.com

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

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 General Operations ASDM Configuration Guide
Copyright © 2012-2014 Cisco Systems, Inc. All rights reserved.



CONTENTS

About This Guide xli

Document Objectives xli

Related Documentation xli

Conventions xlii

Obtaining Documentation and Submitting a Service Request xlii

PART 1 **Getting Started with the ASA**

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

ASDM Requirements 1-1

ASDM Client Operating System and Browser Requirements 1-2

1-2

Java and Browser Compatibility 1-3

Hardware and Software Compatibility 1-7

VPN Compatibility 1-7

New Features 1-7

New Features in ASA 9.1(5)/ASDM 7.1(6) 1-7

New Features in ASA 9.1(4)/ASDM 7.1(5) 1-8

New Features in ASA 9.1(3)/ASDM 7.1(4) 1-10

New Features in ASA 9.1(2)/ASDM 7.1(3) 1-12

New Features in ASA 8.4(6)/ASDM 7.1(2.102) 1-18

New Features in ASA 9.0(2)/ASDM 7.1(2) 1-18

New Features in ASA 9.1(1)/ASDM 7.1(1) 1-18

Firewall Functional Overview 1-18

Security Policy Overview 1-19

Permitting or Denying Traffic with Access Rules 1-19

Applying NAT 1-19

Protecting from IP Fragments 1-19

Using AAA for Through Traffic 1-20

Applying HTTP, HTTPS, or FTP Filtering 1-20

Applying Application Inspection 1-20

Sending Traffic to a Module 1-20

Applying QoS Policies 1-20

Applying Connection Limits and TCP Normalization 1-20

Enabling Threat Detection 1-20
Enabling the Botnet Traffic Filter 1-21
Configuring Cisco Unified Communications 1-21
Firewall Mode Overview 1-21
Stateful Inspection Overview 1-22
VPN Functional Overview 1-23
Security Context Overview 1-23
ASA Clustering Overview 1-24

CHAPTER 2 **Configuring the Switch for Use with the ASA Services Module** 2-1
Information About the Switch 2-1
How the ASA Services Module Works with the Switch 2-1
Supported Switch Hardware and Software 2-3
Backplane Connection 2-4
ASA and IOS Feature Interaction 2-4
Information About SVIs 2-5
Guidelines and Limitations 2-7
Verifying the Module Installation 2-8
Assigning VLANs to the ASA Services Module 2-9
Using the MSFC as a Directly Connected Router (SVIs) 2-12
Configuring the Switch for ASA Failover 2-13
Assigning VLANs to the Secondary ASA Services Module 2-13
Adding a Trunk Between a Primary Switch and Secondary Switch 2-13
Ensuring Compatibility with Transparent Firewall Mode 2-13
Enabling Autostate Messaging for Rapid Link Failure Detection 2-13
Resetting the ASA Services Module 2-14
Monitoring the ASA Services Module 2-14
Feature History for the Switch for Use with the ASA Services Module 2-17

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
Information About Connection Methods 3-3
Logging In 3-4
Logging Out of a Console Session 3-5
Logging Out 3-5
Killing an Active Console Connection 3-6

Logging Out of a Telnet Session 3-6
Configuring ASDM Access for Appliances 3-7
Using the Pre-login Banner in ASDM 3-7
Accessing ASDM Using the Factory Default Configuration 3-7
Customizing ASDM Access (ASA 5505) 3-8
Customizing ASDM Access (ASA 5510 and Higher) 3-10
Configuring ASDM Access for the ASA Services Module 3-12
Starting ASDM 3-14
Connecting to ASDM for the First Time 3-15
Starting ASDM from the ASDM-IDM Launcher 3-16
Starting ASDM from the Java Web Start Application 3-16
Using ASDM in Demo Mode 3-17
Installing an Identity Certificate for ASDM 3-18
Factory Default Configurations 3-19
Restoring the Factory Default Configuration 3-19
ASA 5505 Default Configuration 3-21
ASA 5505 Routed Mode Default Configuration 3-22
ASA 5505 Transparent Mode Sample Configuration 3-23
ASA 5510 and Higher Default Configuration 3-25
Getting Started with the Configuration 3-25
Using the Command Line Interface Tool in ASDM 3-26
Using the Command Line Interface Tool 3-26
Handling Command Errors 3-27
Using Interactive Commands 3-27
Avoiding Conflicts with Other Administrators 3-27
Showing Commands Ignored by ASDM on the Device 3-27
Applying Configuration Changes to Connections 3-28

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 4-11
Status Bar 4-11
Connection to Device 4-12
Device List 4-12
Common Buttons 4-12
Keyboard Shortcuts 4-13
Find Function 4-15
Using the Find Function in Most ASDM Panes 4-15
Using the Find Function in the ACL Manager Pane 4-16
Enabling Extended Screen Reader Support 4-16
Organizational Folder 4-17
About the Help Window 4-17
Home Pane (Single Mode and Context) 4-17
Device Dashboard Tab 4-18
Device Information Pane 4-19
Interface Status Pane 4-20
VPN Sessions Pane 4-20
Failover Status Pane 4-20
System Resources Status Pane 4-20
Traffic Status Pane 4-20
Latest ASDM Syslog Messages Pane 4-21
Firewall Dashboard Tab 4-22
Traffic Overview Pane 4-23
Top 10 Access Rules Pane 4-23
Top Usage Status Pane 4-23
Top Ten Protected Servers Under SYN Attack Pane 4-24
Top 200 Hosts Pane 4-24
Top Botnet Traffic Filter Hits Pane 4-24
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
Time-Based License Activation Guidelines 5-24
How the Time-Based License Timer Works 5-25
How Permanent and Time-Based Licenses Combine 5-25
Stacking Time-Based Licenses 5-26
Time-Based License Expiration 5-26
Shared AnyConnect Premium Licenses 5-27
Information About the Shared Licensing Server and Participants 5-27
Communication Issues Between Participant and Server 5-28
Information About the Shared Licensing Backup Server 5-28
Failover and Shared Licenses 5-29
Maximum Number of Participants 5-30
Failover or ASA Cluster Licenses 5-31
Failover License Requirements and Exceptions 5-31
ASA Cluster License Requirements and Exceptions 5-31
How Failover or ASA Cluster Licenses Combine 5-32
Loss of Communication Between Failover or ASA Cluster Units 5-33
Upgrading Failover Pairs 5-33
No Payload Encryption Models 5-33
Licenses FAQ 5-34
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

Configuring the Shared Licensing Server 5-38
Configuring the Shared Licensing Participant and the Optional Backup Server 5-39
Monitoring Licenses 5-39
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
Information About the Firewall Mode 6-1
Information About Routed Firewall Mode 6-1
Information About Transparent Firewall Mode 6-2
Using the Transparent Firewall in Your Network 6-2
Bridge Groups 6-3
Management Interface (ASA 5510 and Higher) 6-4
Allowing Layer 3 Traffic 6-4
Allowed MAC Addresses 6-5
Passing Traffic Not Allowed in Routed Mode 6-5
Passing Traffic For Routed-Mode Features 6-5
BPDU Handling 6-5
MAC Address vs. Route Lookups 6-6
ARP Inspection 6-6
MAC Address Table 6-7
Licensing Requirements for the Firewall Mode 6-7
Default Settings 6-7
Guidelines and Limitations 6-8
Setting the Firewall Mode (Single Mode) 6-9
Configuring ARP Inspection for the Transparent Firewall 6-10
Task Flow for Configuring ARP Inspection 6-10
Adding a Static ARP Entry 6-10
Enabling ARP Inspection 6-11
Customizing the MAC Address Table for the Transparent Firewall 6-12
Adding a Static MAC Address 6-12
Disabling MAC Address Learning 6-13
Firewall Mode Examples 6-13
How Data Moves Through the ASA in Routed Firewall Mode 6-13
An Inside User Visits a Web Server 6-14
An Outside User Visits a Web Server on the DMZ 6-15
An Inside User Visits a Web Server on the DMZ 6-16

An Outside User Attempts to Access an Inside Host 6-16
A DMZ User Attempts to Access an Inside Host 6-18
How Data Moves Through the Transparent Firewall 6-19
An Inside User Visits a Web Server 6-20
An Inside User Visits a Web Server Using NAT 6-21
An Outside User Visits a Web Server on the Inside Network 6-22
An Outside User Attempts to Access an Inside Host 6-23
Feature History for the Firewall Mode 6-24

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
Interface Selection (ASA 5505) 7-3
Switch Port Allocation (ASA 5505) 7-3
Interface IP Address Configuration (ASA 5505, Routed Mode) 7-3
Interface Configuration - PPPoE (ASA 5505, Routed Mode, Single Mode) 7-3
Outside Interface Configuration (ASA 5510 and Higher, Routed Mode) 7-4
Outside Interface Configuration - PPPoE (ASA 5510 and Higher, Routed Mode, Single Mode) 7-4
Management IP Address Configuration (Transparent Mode) 7-4
Other Interfaces Configuration (ASA 5510 and Higher) 7-4
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
ASA CX Basic Configuration (ASA 5585-X) 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 2 **Configuring High Availability and Scalability**

CHAPTER 8 **Configuring Multiple Context Mode** 8-1

Information About Security Contexts 8-1
Common Uses for Security Contexts 8-2
Context Configuration Files 8-2
Context Configurations 8-2
System Configuration 8-2
Admin Context Configuration 8-2
How the ASA Classifies Packets 8-3
Valid Classifier Criteria 8-3
Classification Examples 8-4
Cascading Security Contexts 8-6
Management Access to Security Contexts 8-7
System Administrator Access 8-7
Context Administrator Access 8-8
Information About Resource Management 8-8
Resource Classes 8-8
Resource Limits 8-8
Default Class 8-9
Using Oversubscribed Resources 8-10
Using Unlimited Resources 8-11
Information About MAC Addresses 8-11
Default MAC Address 8-12
Interaction with Manual MAC Addresses 8-12
Failover MAC Addresses 8-12
MAC Address Format 8-12
Licensing Requirements for Multiple Context Mode 8-13
Prerequisites 8-14
Guidelines and Limitations 8-14
Default Settings 8-15
Configuring Multiple Contexts 8-15
Task Flow for Configuring Multiple Context Mode 8-15
Enabling or Disabling Multiple Context Mode 8-16
Enabling Multiple Context Mode 8-16
Restoring Single Context Mode 8-16
Configuring a Class for Resource Management 8-17
Configuring a Security Context 8-20
Automatically Assigning MAC Addresses to Context Interfaces 8-24
Changing Between Contexts and the System Execution Space 8-25
Managing Security Contexts 8-26

Removing a Security Context 8-26
Changing the Admin Context 8-27
Changing the Security Context URL 8-28
Reloading a Security Context 8-29
Reloading by Clearing the Configuration 8-30
Reloading by Removing and Re-adding the Context 8-31
Monitoring Security Contexts 8-31
Monitoring Context Resource Usage 8-31
Viewing Assigned MAC Addresses 8-32
Viewing MAC Addresses in the System Configuration 8-33
Viewing MAC Addresses Within a Context 8-33
Feature History for Multiple Context Mode 8-34

CHAPTER 9 **Configuring Failover** 9-1
Introduction to Failover 9-1
Failover Overview 9-2
Failover System Requirements 9-2
Hardware Requirements 9-2
Software Requirements 9-2
License Requirements 9-3
Failover and Stateful Failover Links 9-3
Failover Link 9-3
Stateful Failover Link 9-4
Avoiding Interrupted Failover and Data Links 9-5
MAC Addresses and IP Addresses 9-7
Intra- and Inter-Chassis Module Placement for the ASA Services Module 9-8
Intra-Chassis Failover 9-8
Inter-Chassis Failover 9-9
Stateless and Stateful Failover 9-12
Stateless Failover 9-13
Stateful Failover 9-13
Transparent Firewall Mode Requirements 9-14
Transparent Mode Requirements for Appliances 9-14
Transparent Mode Requirements for Modules 9-15
Failover Health Monitoring 9-16
Unit Health Monitoring 9-16
Interface Monitoring 9-16
Failover Times 9-17
Configuration Synchronization 9-17

Running Configuration Replication 9-18
Command Replication 9-18
Information About Active/Standby Failover 9-19
Primary/Secondary Roles and Active/Standby Status 9-19
Active Unit Determination at Startup 9-19
Failover Events 9-20
Information About Active/Active Failover 9-20
Active/Active Failover Overview 9-21
Primary/Secondary Roles and Active/Standby Status for a Failover Group 9-21
Active Unit Determination for Failover Groups at Startup 9-21
Failover Events 9-22
Licensing Requirements Failover 9-23
Prerequisites for Failover 9-23
Guidelines and Limitations 9-24
Default Settings 9-24
Configuring Active/Standby Failover 9-25
Configuring Active/Active Failover 9-32
Configuring Optional Failover Parameters 9-41
Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses 9-41
Configuring Interface Monitoring and Standby Addresses 9-44
Configuring Support for Asymmetrically Routed Packets (Active/Active Mode) 9-45
Managing Failover 9-47
Modifying the Failover Setup 9-47
Forcing Failover 9-50
Disabling Failover 9-51
Restoring a Failed Unit 9-51
Re-Syncing the Configuration 9-52
Monitoring Failover 9-52
Failover Messages 9-52
Failover Syslog Messages 9-52
Failover Debug Messages 9-53
SNMP Failover Traps 9-53
Monitoring Failover 9-53
System 9-53
Failover Group 1 and Failover Group 2 9-54
Feature History for Failover 9-54

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 Hardware and Software Requirements 10-3
Bootstrap Configuration 10-3
Master and Slave Unit Roles 10-3
Master Unit Election 10-3
ASA Cluster Interfaces 10-4
Interface Types 10-4
Interface Type Mode 10-6
Cluster Control Link 10-6
Cluster Control Link Traffic Overview 10-7
Cluster Control Link Interfaces and Network 10-7
Sizing the Cluster Control Link 10-7
Cluster Control Link Redundancy 10-8
Cluster Control Link Reliability 10-8
Cluster Control Link Failure 10-9
High Availability within the ASA Cluster 10-9
Unit Health Monitoring 10-9
Interface monitoring 10-9
Unit or Interface Failure 10-9
Rejoining the Cluster 10-10
Data Path Connection State Replication 10-10
Configuration Replication 10-11
ASA Cluster Management 10-11
Management Network 10-11
Management Interface 10-11
Master Unit Management Vs. Slave Unit Management 10-12
RSA Key Replication 10-12
ASDM Connection Certificate IP Address Mismatch 10-12
Load Balancing Methods 10-13
Spanned EtherChannel (Recommended) 10-13
Policy-Based Routing (Routed Firewall Mode Only) 10-15
Equal-Cost Multi-Path Routing (Routed Firewall Mode Only) 10-16
Inter-Site Clustering 10-16
Inter-Site Clustering Guidelines 10-16
Sizing the Data Center Interconnect 10-17

Inter-Site Example 10-17
How the ASA Cluster Manages Connections 10-18
Connection Roles 10-18
New Connection Ownership 10-19
Sample Data Flow 10-19
Rebalancing New TCP Connections Across the Cluster 10-20
ASA Features and Clustering 10-20
Unsupported Features 10-20
Centralized Features 10-21
Features Applied to Individual Units 10-22
Dynamic Routing 10-22
Multicast Routing 10-24
NAT 10-24
AAA for Network Access 10-25
Syslog and Netflow 10-26
SNMP 10-26
VPN 10-26
FTP 10-26
Cisco TrustSec 10-27
Licensing Requirements for ASA Clustering 10-27
Prerequisites for ASA Clustering 10-27
Guidelines and Limitations 10-28
Default Settings 10-32
Configuring ASA Clustering 10-32
Task Flow for ASA Cluster Configuration 10-32
Cabling the Cluster Units and Configuring Upstream and Downstream Equipment 10-33
Backing Up Your Configurations (Recommended) 10-35
Configuring the Cluster Interface Mode on the Master Unit 10-35
(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit 10-38
Configuring Individual Interfaces (Recommended for the Management Interface) 10-38
Configuring Spanned EtherChannels 10-43
Adding or Joining an ASA Cluster 10-50
Managing ASA Cluster Members 10-59
Configuring ASA Cluster Parameters 10-59
Adding a New Slave from the Master Unit 10-62
Becoming an Inactive Member 10-64
Inactivating a Slave Member from the Master Unit 10-65
Leaving the Cluster 10-66

Changing the Master Unit 10-67
Executing a Command Cluster-Wide 10-69
Monitoring the ASA Cluster 10-70
Cluster Dashboards 10-70
Monitoring Screens 10-70
Viewing the Cluster Summary 10-70
Monitoring Cluster Resources 10-71
Monitoring Cluster Traffic 10-71
Monitoring the Cluster Control Link 10-72
Related Features 10-72
Configuration Examples for ASA Clustering 10-72
Sample ASA and Switch Configuration 10-72
ASA Configuration 10-73
IOS Switch Configuration 10-74
Firewall on a Stick 10-75
Traffic Segregation 10-77
Redundant Interface (PBR or ECMP) 10-80
Spanned EtherChannel With Backup Links 10-82
Feature History for ASA Clustering 10-88

PART 3 **Configuring Interfaces**

CHAPTER 11 **Starting Interface Configuration (ASA 5510 and Higher) 11-1**
Information About Starting ASA 5510 and Higher Interface Configuration 11-2
Auto-MDI/MDIX Feature 11-2
Interfaces in Transparent Mode 11-2
Management Interface 11-2
Management Interface Overview 11-2
Management *Slot*/*Port* Interface 11-3
Using Any Interface for Management-Only Traffic 11-3
Management Interface for Transparent Mode 11-4
No Support for Redundant Management Interfaces 11-4
Management 0/0 Interface on the ASA 5512-X through ASA 5555-X 11-4
Redundant Interfaces 11-5
Redundant Interface MAC Address 11-5
EtherChannels 11-5
Channel Group Interfaces 11-5
Connecting to an EtherChannel on Another Device 11-5
Link Aggregation Control Protocol 11-6

Load Balancing 11-7
EtherChannel MAC Address 11-8
Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size 11-8
MTU Overview 11-8
Default MTU 11-8
Path MTU Discovery 11-9
Setting the MTU and Jumbo Frames 11-9
TCP Maximum Segment Size Overview 11-9
Default TCP MSS 11-9
Setting the TCP MSS for VPN and Non-VPN Traffic 11-9
Licensing Requirements for ASA 5510 and Higher Interfaces 11-10
Guidelines and Limitations 11-12
Default Settings 11-14
Starting Interface Configuration (ASA 5510 and Higher) 11-14
Task Flow for Starting Interface Configuration 11-15
Converting In-Use Interfaces to a Redundant or EtherChannel Interface 11-16
Enabling the Physical Interface and Configuring Ethernet Parameters 11-25
Configuring a Redundant Interface 11-27
Configuring a Redundant Interface 11-27
Changing the Active Interface 11-30
Configuring an EtherChannel 11-30
Adding Interfaces to the EtherChannel 11-31
Customizing the EtherChannel 11-33
Configuring VLAN Subinterfaces and 802.1Q Trunking 11-36
Enabling Jumbo Frame Support (Supported Models) 11-39
Monitoring Interfaces 11-40
ARP Table 11-40
MAC Address Table 11-40
Interface Graphs 11-40
Graph/Table 11-42
Where to Go Next 11-43
Feature History for ASA 5510 and Higher Interfaces 11-44

CHAPTER 12 **Starting Interface Configuration (ASA 5505)** 12-1

Information About ASA 5505 Interfaces 12-1
Understanding ASA 5505 Ports and Interfaces 12-2
Maximum Active VLAN Interfaces for Your License 12-2
VLAN MAC Addresses 12-4

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

CHAPTER 13 **Completing Interface Configuration (Routed Mode)** 13-1
Information About Completing Interface Configuration in Routed Mode 13-1
Security Levels 13-1
Dual IP Stack (IPv4 and IPv6) 13-2
Licensing Requirements for Completing Interface Configuration in Routed Mode 13-2
Guidelines and Limitations 13-5
Default Settings 13-6
Completing Interface Configuration in Routed Mode 13-6
Task Flow for Completing Interface Configuration 13-7
Configuring General Interface Parameters 13-7
PPPoE IP Address and Route Settings 13-11
Configuring the MAC Address, MTU, and TCP MSS 13-13
Configuring IPv6 Addressing 13-15
Information About IPv6 13-15
Configuring a Global IPv6 Address 13-16
Configuring IPv6 Neighbor Discovery 13-18
(Optional) Configuring the Link-Local Addresses Automatically 13-18
(Optional) Configuring the Link-Local Addresses Manually 13-19
Allowing Same Security Level Communication 13-20

Turning Off and Turning On Interfaces 13-22
Monitoring Interfaces 13-22
ARP Table 13-23
DHCP 13-23
DHCP Server Table 13-23
DHCP Client Lease Information 13-24
DHCP Statistics 13-25
MAC Address Table 13-26
Dynamic ACLs 13-26
Interface Graphs 13-26
Graph/Table 13-28
PPPoE Client 13-29
Interface Connection 13-29
Track Status for 13-29
Monitoring Statistics for 13-29
Feature History for Interfaces in Routed Mode 13-30

CHAPTER 14 **Completing Interface Configuration (Transparent Mode, 8.4 and Later)** 14-1
Information About Completing Interface Configuration in Transparent Mode (8.4 and Later) 14-1
Bridge Groups in Transparent Mode 14-2
Security Levels 14-2
Licensing Requirements for Completing Interface Configuration in Transparent Mode 14-3
Guidelines and Limitations 14-5
Default Settings 14-7
Completing Interface Configuration in Transparent Mode (8.4 and Later) 14-7
Task Flow for Completing Interface Configuration 14-8
Configuring Bridge Groups 14-8
Configuring General Interface Parameters 14-9
Configuring a Management Interface (ASA 5510 and Higher) 14-12
Configuring the MAC Address, MTU, and TCP MSS 14-15
Configuring IPv6 Addressing 14-17
Information About IPv6 14-17
Configuring a Global IPv6 Address 14-18
Configuring IPv6 Neighbor Discovery 14-20
(Optional) Configuring the Link-Local Addresses Automatically 14-20
(Optional) Configuring the Link-Local Addresses Manually 14-21
Allowing Same Security Level Communication 14-21
Turning Off and Turning On Interfaces 14-22
Monitoring Interfaces 14-22

ARP Table 14-23
DHCP 14-23
DHCP Server Table 14-23
DHCP Client Lease Information 14-24
DHCP Statistics 14-25
MAC Address Table 14-26
Dynamic ACLs 14-26
Interface Graphs 14-26
Graph/Table 14-28
PPPoE Client 14-29
Interface Connection 14-29
Track Status for 14-29
Monitoring Statistics for 14-29
Feature History for Interfaces in Transparent Mode 14-30

CHAPTER 15 **Completing Interface Configuration (Transparent Mode, 8.3 and Earlier)** 15-1
Information About Completing Interface Configuration in Transparent Mode (8.3 and Earlier) 15-1
Information About the Global Management IP Address 15-2
Security Levels 15-2
Licensing Requirements for Completing Interface Configuration in Transparent Mode 15-3
Guidelines and Limitations 15-3
Default Settings 15-4
Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier) 15-4
Configuring the IPv4 Address 15-4
Configuring the IPv6 Address 15-5
Information About IPv6 15-5
Configuring the Global Address 15-7
Configuring the Link-Local Addresses Automatically 15-7
Configuring the Link-Local Address on an Interface Manually 15-8
Configuring DAD Settings 15-8
Completing Interface Configuration in Transparent Mode (8.3 and Earlier) 15-9
Task Flow for Completing Interface Configuration 15-9
Configuring General Interface Parameters 15-9
Configuring a Management Interface (ASA 5510 and Higher) 15-11
Configuring General Parameters and the IPv4 Address 15-11
Configuring a Global IPv6 Address and Other Options 15-13
Configuring the MAC Address, MTU, and TCP MSS 15-14
Allowing Same Security Level Communication 15-17
Monitoring Interfaces 15-17

Feature History for Interfaces in Transparent Mode 15-18

PART 4 Configuring Basic Settings

CHAPTER 16 Configuring Basic Settings 16-1

Configuring the Hostname, Domain Name, and Passwords 16-1
Setting the Hostname, Domain Name, and the enable and Telnet Passwords 16-1
Feature History for the Hostname, Domain Name, and Passwords 16-3
Setting the Date and Time 16-3
Setting the Date and Time Using an NTP Server 16-3
Adding or Editing the NTP Server Configuration 16-4
Setting the Date and Time Manually 16-4
Configuring the Master Passphrase 16-5
Information About the Master Passphrase 16-5
Licensing Requirements for the Master Passphrase 16-6
Guidelines and Limitations 16-6
Adding or Changing the Master Passphrase 16-6
Disabling the Master Passphrase 16-7
Recovering the Master Passphrase 16-8
Feature History for the Master Passphrase 16-8
Configuring the DNS Server 16-8
Changing the Heap Memory Size 16-10
Monitoring DNS Cache 16-10

CHAPTER 17 Configuring DHCP Services 17-1

Information About DHCP Services 17-1
Information About the DHCP Server 17-1
Information About the DHCP Relay Agent 17-2
Licensing Requirements for DHCP 17-2
Guidelines and Limitations 17-2
Configuring DHCP Services 17-4
Configuring the DHCP Server 17-4
Enabling the DHCP Server 17-5
Configuring Advanced DHCP Options 17-6
Configuring the DHCP Relay Agent 17-7
Additional References 17-9
RFCs 17-9
Monitoring DHCP Services 17-9

Feature History for DHCP Services 17-10

CHAPTER 18 **Configuring Dynamic DNS** 18-1

Information About DDNS 18-1

Licensing Requirements for DDNS 18-2

Guidelines and Limitations 18-2

Configuring Dynamic DNS 18-3

DDNS Monitoring 18-7

Feature History for DDNS 18-7

CHAPTER 19 **Configuring Web Cache Services Using WCCP** 19-1

Information About WCCP 19-1

Guidelines and Limitations 19-1

Licensing Requirements for WCCP 19-2

Configuring WCCP Service Groups 19-3

Adding or Editing WCCP Service Groups 19-3

Configuring Packet Redirection 19-4

Adding or Editing Packet Redirection 19-4

WCCP Monitoring 19-4

Feature History for WCCP 19-5

PART 5 **Configuring Objects and ACLs**

CHAPTER 20 **Configuring Objects** 20-1

Information About Objects 20-1

Licensing Requirements for Objects 20-1

Guidelines and Limitations 20-1

Configuring Objects 20-2

Configuring Network Objects and Groups 20-2

Configuring a Network Object 20-3

Configuring a Network Object Group 20-3

Configuring Service Objects and Service Groups 20-4

Configuring a Service Object 20-5

Configuring a Service Group 20-5

Configuring a TCP or UDP Port Service Group 20-6

Configuring an ICMP Group 20-6

Configuring a Protocol Group 20-7

Configuring Local User Groups 20-7
Configuring Security Group Object Groups 20-8
Configuring Regular Expressions 20-10
Creating a Regular Expression 20-10
Creating a Regular Expression Class Map 20-14
Configuring Time Ranges 20-15
Monitoring Objects 20-16
Feature History for Objects 20-17

CHAPTER 21 **Using the ACL Manager** 21-1
Information About the ACL Manager 21-1
Licensing Requirements for the ACL Manager 21-1
Guidelines and Limitations 21-2
Adding ACLs and ACEs 21-2
Using Standard ACLs in the ACL Manager 21-4
Feature History for the ACL Manager 21-5

CHAPTER 22 **Adding a Standard Access Control List** 22-1
Information About Standard ACLs 22-1
Licensing Requirements for Standard ACLs 22-1
Guidelines and Limitations 22-1
Default Settings 22-2
Adding Standard ACLs 22-3
Using Standard ACLs 22-3
Adding a Standard ACL 22-3
Adding an ACE to a Standard ACL 22-3
Editing an ACE in a Standard ACL 22-4
Feature History for Standard ACLs 22-4

CHAPTER 23 **Adding a Webtype Access Control List** 23-1
Licensing Requirements for Webtype ACLs 23-1
Guidelines and Limitations 23-1
Default Settings 23-3
Using Webtype ACLs 23-3
Task Flow for Configuring Webtype ACLs 23-3
Adding a Webtype ACL and ACE 23-3
Editing Webtype ACLs and ACEs 23-5

Deleting Webtype ACLs and ACEs 23-5
Feature History for Webtype ACLs 23-6

PART 6 Configuring IP Routing

CHAPTER 24 Routing Overview 24-1
Information About Routing 24-1
Switching 24-1
Path Determination 24-2
Supported Route Types 24-2
Static Versus Dynamic 24-3
Single-Path Versus Multipath 24-3
Flat Versus Hierarchical 24-3
Link-State Versus Distance Vector 24-3
How Routing Behaves Within the ASA 24-4
Egress Interface Selection Process 24-4
Next Hop Selection Process 24-4
Supported Internet Protocols for Routing 24-5
Information About the Routing Table 24-5
Displaying the Routing Table 24-6
How the Routing Table Is Populated 24-6
Administrative Distances for Routes 24-7
Backup Routes 24-8
How Forwarding Decisions Are Made 24-8
Dynamic Routing and Failover 24-9
Dynamic Routing and Clustering 24-9
Dynamic Routing in Multiple Context Mode 24-10
Route Resource Management 24-11
Disabling Proxy ARP Requests 24-11

CHAPTER 25 Configuring Static and Default Routes 25-1
Information About Static and Default Routes 25-1
Licensing Requirements for Static and Default Routes 25-2
Guidelines and Limitations 25-2
Configuring Static and Default Routes 25-2
Configuring a Static Route 25-3
Adding or Editing a Static Route 25-3
Configuring Static Route Tracking 25-6

Deleting Static Routes 25-6
Configuring a Default Static Route 25-7
Limitations on Configuring a Default Static Route 25-7
Configuring IPv6 Default and Static Routes 25-8
Monitoring a Static or Default Route 25-9
Configuration Examples for Static or Default Routes 25-11
Feature History for Static and Default Routes 25-13

CHAPTER 26 **Defining Route Maps** 26-1
Information About Route Maps 26-1
Permit and Deny Clauses 26-2
Match and Set Clause Values 26-2
Licensing Requirements for Route Maps 26-3
Guidelines and Limitations 26-3
Defining a Route Map 26-4
Adding or Editing a Route Map 26-4
Customizing a Route Map 26-5
Defining a Route to Match a Specific Destination Address 26-6
Configuring Prefix Lists 26-7
Configuring Prefix Rules 26-8
Configuring the Metric Values for a Route Action 26-8
Configuration Example for Route Maps 26-9
Feature History for Route Maps 26-10

CHAPTER 27 **Configuring OSPF** 27-1
Information About OSPF 27-1
Implementation Differences Between OSPFv2 and OSPFv3 27-3
Using Clustering 27-3
Licensing Requirements for OSPF 27-3
Guidelines and Limitations 27-3
Configuring OSPFv2 27-5
Customizing OSPFv2 27-7
Redistributing Routes Into OSPFv2 27-7
Configuring Route Summarization When Redistributing Routes Into OSPFv2 27-11
Adding a Route Summary Address 27-11
Adding or Editing an OSPF Summary Address 27-12
Configuring Route Summarization Between OSPFv2 Areas 27-13
Configuring OSPFv2 Interface Parameters 27-14

Configuring OSPFv2 Area Parameters 27-19
Configuring an OSPFv2 NSSA 27-20
Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3) 27-22
Defining Static OSPFv2 Neighbors 27-24
Configuring Route Calculation Timers 27-25
Logging Neighbors Going Up or Down 27-26
Configuring Filtering in OSPF 27-27
Configuring a Virtual Link in OSPF 27-28
Configuring OSPFv3 27-29
Enabling OSPFv3 27-30
Configuring OSPFv3 Interface Parameters 27-31
Configuring OSPFv3 Router Parameters 27-38
Configuring OSPFv3 Area Parameters 27-40
Configuring a Virtual Link Neighbor 27-42
Configuring OSPFv3 Passive Interfaces 27-44
Configuring OSPFv3 Administrative Distance 27-45
Configuring OSPFv3 Timers 27-46
Defining Static OSPFv3 Neighbors 27-50
Resetting OSPFv3 Default Parameters 27-52
Sending Syslog Messages 27-53
Suppressing Syslog Messages 27-54
Calculating Summary Route Costs 27-54
Generating a Default External Route into an OSPFv3 Routing Domain 27-55
Configuring an IPv6 Summary Prefix 27-56
Redistributing IPv6 Routes 27-57
Removing the OSPF Configuration 27-60
Configuration Example for OSPFv2 27-60
Configuration Examples for OSPFv3 27-63
Monitoring OSPF 27-65
Additional References 27-69
RFCs 27-69
Feature History for OSPF 27-69

CHAPTER 28 **Configuring EIGRP** 28-1
Information About EIGRP 28-1
Using Clustering 28-2
Licensing Requirements for EIGRP 28-2
Guidelines and Limitations 28-3

Task List to Configure an EIGRP Process 28-3
Configuring EIGRP 28-4
Enabling EIGRP 28-4
Enabling EIGRP Stub Routing 28-6
Customizing EIGRP 28-7
Defining a Network for an EIGRP Routing Process 28-8
Configuring Interfaces for EIGRP 28-9
Configuring Passive Interfaces 28-12
Configuring the Summary Aggregate Addresses on Interfaces 28-13
Changing the Interface Delay Value 28-14
Enabling EIGRP Authentication on an Interface 28-15
Defining an EIGRP Neighbor 28-17
Redistributing Routes Into EIGRP 28-18
Filtering Networks in EIGRP 28-21
Customizing the EIGRP Hello Interval and Hold Time 28-23
Disabling Automatic Route Summarization 28-24
Configuring Default Information in EIGRP 28-25
Disabling EIGRP Split Horizon 28-26
Restarting the EIGRP Process 28-27
Monitoring EIGRP 28-28
Configuration Example for EIGRP 28-29
Feature History for EIGRP 28-30

CHAPTER 29 **Configuring RIP** 29-1

Information About RIP 29-1
Routing Update Process 29-2
RIP Routing Metric 29-2
RIP Stability Features 29-2
RIP Timers 29-2
Using Clustering 29-3
Licensing Requirements for RIP 29-3
Guidelines and Limitations 29-3
Configuring RIP 29-4
Enabling RIP 29-4
Customizing RIP 29-5
Configuring the RIP Version 29-5
Configuring Interfaces for RIP 29-7
Editing a RIP Interface 29-8

Configuring the RIP Send and Receive Version on an Interface 29-9
Configuring Route Summarization 29-10
Filtering Networks in RIP 29-11
Adding or Editing a Filter Rule 29-12
Redistributing Routes into the RIP Routing Process 29-13
Enabling RIP Authentication 29-15
Restarting the RIP Process 29-16
Monitoring RIP 29-16
Configuration Example for RIP 29-17
Feature History for RIP 29-18

CHAPTER 30 **Configuring Multicast Routing** 30-1
Information About Multicast Routing 30-1
Stub Multicast Routing 30-2
PIM Multicast Routing 30-2
Multicast Group Concept 30-2
Multicast Addresses 30-3
Clustering 30-3
Licensing Requirements for Multicast Routing 30-3
Guidelines and Limitations 30-3
Enabling Multicast Routing 30-4
Customizing Multicast Routing 30-4
Configuring Stub Multicast Routing and Forwarding IGMP Messages 30-5
Configuring a Static Multicast Route 30-6
Configuring IGMP Features 30-7
Disabling IGMP on an Interface 30-8
Configuring IGMP Group Membership 30-8
Configuring a Statically Joined IGMP Group 30-9
Controlling Access to Multicast Groups 30-10
Limiting the Number of IGMP States on an Interface 30-11
Modifying the Query Messages to Multicast Groups 30-12
Changing the IGMP Version 30-13
Configuring PIM Features 30-14
Enabling and Disabling PIM on an Interface 30-14
Configuring a Static Rendezvous Point Address 30-15
Configuring the Designated Router Priority 30-16
Configuring and Filtering PIM Register Messages 30-17
Configuring PIM Message Intervals 30-18
Configuring a Route Tree 30-18

Configuring a Multicast Group 30-19
Filtering PIM Neighbors 30-19
Configuring a Bidirectional Neighbor Filter 30-21
Configuring a Multicast Boundary 30-22
Configuration Example for Multicast Routing 30-23
Additional References 30-25
Related Documents 30-25
RFCs 30-25
Feature History for Multicast Routing 30-25

CHAPTER 31 **Configuring IPv6 Neighbor Discovery** 31-1
Information About IPv6 Neighbor Discovery 31-1
Neighbor Solicitation Messages 31-2
Neighbor Reachable Time 31-3
Duplicate Address Detection 31-3
Router Advertisement Messages 31-3
Static IPv6 Neighbors 31-5
Licensing Requirements for IPv6 Neighbor Discovery 31-5
Prerequisites for IPv6 Neighbor Discovery 31-5
Guidelines and Limitations 31-5
Default Settings for IPv6 Neighbor Discovery 31-7
Configuring IPv6 Neighbor Discovery 31-7
Entering Interface Configuration Mode 31-7
Configuring the Neighbor Solicitation Message Interval 31-8
Configuring the Neighbor Reachable Time 31-9
Configuring the Router Advertisement Transmission Interval 31-10
Configuring the Router Lifetime Value 31-11
Configuring DAD Settings 31-12
Suppressing Router Advertisement Messages 31-13
Configuring Address Config Flags for IPv6 DHCP Relay 31-14
Configuring the IPv6 Prefix in Router Advertisements 31-15
Configuring a Static IPv6 Neighbor 31-17
Viewing and Clearing Dynamically Discovered Neighbors 31-18
Monitoring IPv6 Neighbor Discovery 31-18
Additional References 31-18
Related Documents for IPv6 Prefixes 31-20
RFCs for IPv6 Prefixes and Documentation 31-20
Feature History for IPv6 Neighbor Discovery 31-20

PART 7 Configuring AAA Servers and the Local Database

CHAPTER 32 Information About AAA 32-1

Authentication 32-1
Authorization 32-2
Accounting 32-2
Interaction Between Authentication, Authorization, and Accounting 32-2
AAA Servers 32-2
AAA Server Groups 32-3
Local Database Support 32-3
Summary of AAA Service Support 32-3

CHAPTER 33 Configuring the Local Database for AAA 33-1

Information About the Local Database 33-1
Fallback Support 33-2
How Fallback Works with Multiple Servers in a Group 33-2
Licensing Requirements for the Local Database 33-3
Guidelines and Limitations 33-3
Adding a User Account to the Local Database 33-3
Testing Local Database Authentication and Authorization 33-7
Monitoring the Local Database 33-7
Feature History for the Local Database 33-8

CHAPTER 34 Configuring RADIUS Servers for AAA 34-1

Information About RADIUS Servers 34-1
Supported Authentication Methods 34-2
User Authorization of VPN Connections 34-2
Supported Sets of RADIUS Attributes 34-2
Supported RADIUS Authorization Attributes 34-3
Supported IETF RADIUS Authorization Attributes 34-12
RADIUS Accounting Disconnect Reason Codes 34-13
Licensing Requirements for RADIUS Servers 34-13
Guidelines and Limitations 34-14
Configuring RADIUS Servers 34-14
Task Flow for Configuring RADIUS Servers 34-14
Configuring RADIUS Server Groups 34-15
Adding a RADIUS Server to a Group 34-16

Adding an Authentication Prompt 34-18
Testing RADIUS Server Authentication and Authorization 34-18
Monitoring RADIUS Servers 34-19
Additional References 34-20
RFCs 34-20
Feature History for RADIUS Servers 34-20

CHAPTER 35 **Configuring TACACS+ Servers for AAA** 35-1
Information About TACACS+ Servers 35-1
Using TACACS+ Attributes 35-1
Licensing Requirements for TACACS+ Servers 35-2
Guidelines and Limitations 35-3
Configuring TACACS+ Servers 35-3
Task Flow for Configuring TACACS+ Servers 35-3
Configuring TACACS+ Server Groups 35-4
Adding a TACACS+ Server to a Group 35-4
Adding an Authentication Prompt 35-5
Testing TACACS+ Server Authentication and Authorization 35-6
Monitoring TACACS+ Servers 35-7
Feature History for TACACS+ Servers 35-7

CHAPTER 36 **Configuring LDAP Servers for AAA** 36-1
Information About LDAP and the ASA 36-1
LDAP Server Guidelines 36-1
How Authentication Works with LDAP 36-2
About the LDAP Hierarchy 36-2
Searching the LDAP Hierarchy 36-3
About Binding to an LDAP Server 36-4
Licensing Requirements for LDAP Servers 36-4
Guidelines and Limitations 36-4
Configuring LDAP Servers 36-5
Task Flow for Configuring LDAP Servers 36-5
Configuring LDAP Attribute Maps 36-5
Configuring LDAP Server Groups 36-7
Adding an LDAP Server to a Group 36-8
Testing LDAP Server Authentication and Authorization 36-9
Monitoring LDAP Servers 36-10
Feature History for LDAP Servers 36-10

CHAPTER 37 **Configuring Windows NT Servers for AAA** 37-1

Information About Windows NT Servers 37-1
Licensing Requirements for Windows NT Servers 37-1
Guidelines and Limitations 37-2
Configuring Windows NT Servers 37-2
Task Flow for Configuring Windows NT Servers 37-2
Configuring Windows NT Server Groups 37-3
Adding a Windows NT Server to a Group 37-3
Testing Windows NT Server Authentication and Authorization 37-4
Monitoring Windows NT Servers 37-5
Feature History for Windows NT Servers 37-5

CHAPTER 38 **Configuring the Identity Firewall** 38-1

Information About the Identity Firewall 38-1
Overview of the Identity Firewall 38-1
Architecture for Identity Firewall Deployments 38-2
Features of the Identity Firewall 38-3
Deployment Scenarios 38-4
Licensing for the Identity Firewall 38-7
Guidelines and Limitations 38-8
Prerequisites 38-9
Configuring the Identity Firewall 38-10
Task Flow for Configuring the Identity Firewall 38-10
Configuring the Active Directory Domain 38-11
Configuring Active Directory Server Groups 38-12
Configuring Active Directory Agents 38-12
Configuring Active Directory Agent Groups 38-13
Configuring Identity Options 38-13
Configuring Identity-Based Security Policy 38-16
Monitoring the Identity Firewall 38-17
Monitoring AD Agents 38-17
Monitoring Groups 38-17
Monitoring Memory Usage for the Identity Firewall 38-18
Monitoring Users for the Identity Firewall 38-18
Feature History for the Identity Firewall 38-19

CHAPTER 39 **Configuring the ASA to Integrate with Cisco TrustSec** 39-1

Information About the ASA Integrated with Cisco TrustSec 39-1
Information about Cisco TrustSec 39-2
About SGT and SXP Support in Cisco TrustSec 39-2
Roles in the Cisco TrustSec Feature 39-3
Security Group Policy Enforcement 39-4
How the ASA Enforces Security Group-Based Policies 39-4
Effects of Changes to Security Groups on the ISE 39-6
About Speaker and Listener Roles on the ASA 39-6
SXP Chattiness 39-7
SXP Timers 39-8
IP-SGT Manager Database 39-8
Features of the ASA-Cisco TrustSec Integration 39-9
Licensing Requirements for Cisco TrustSec 39-11
Prerequisites for Using Cisco TrustSec 39-11
Registering the ASA with the ISE 39-11
Creating a Security Group on the ISE 39-12
Generating the PAC File 39-12
Guidelines and Limitations 39-12
Configuring the ASA for Cisco TrustSec Integration 39-14
Task Flow for Configuring the ASA to Integrate with Cisco TrustSec 39-15
Configuring the AAA Server for Cisco TrustSec Integration 39-15
Importing a Protected Access Credential (PAC) File 39-16
Configuring the Security Exchange Protocol (SXP) 39-17
Adding an SXP Connection Peer 39-19
Refreshing Environment Data 39-20
Configuring the Security Policy 39-20
Monitoring Cisco TrustSec 39-21
Feature History for the Cisco TrustSec Integration 39-22

CHAPTER 40 **Configuring Digital Certificates** 40-1

Information About Digital Certificates 40-1
Public Key Cryptography 40-3
Certificate Scalability 40-3
Key Pairs 40-4
Trustpoints 40-4
Certificate Enrollment 40-4
Proxy for SCEP Requests 40-5

Revocation Checking 40-5
Supported CA Servers 40-5
CRLs 40-6
OCSP 40-6
The Local CA 40-7
Storage for Local CA Files 40-7
The Local CA Server 40-8
Using Certificates and User Login Credentials 40-8
Using User Login Credentials 40-8
Using Certificates 40-9
Licensing Requirements for Digital Certificates 40-9
Prerequisites for Local Certificates 40-10
Prerequisites for SCEP Proxy Support 40-10
Guidelines and Limitations 40-10
Configuring Digital Certificates 40-11
Configuring CA Certificate Authentication 40-12
Adding or Installing a CA Certificate 40-13
Editing or Removing a CA Certificate Configuration 40-14
Showing CA Certificate Details 40-14
Configuring CA Certificate Authentication 40-14
Adding or Installing a CA Certificate 40-15
Editing or Removing a CA Certificate Configuration 40-15
Showing CA Certificate Details 40-16
Configuring CA Certificates for Revocation 40-16
Configuring CRL Retrieval Policy 40-17
Configuring CRL Retrieval Methods 40-17
Configuring OCSP Rules 40-18
Configuring Advanced CRL and OCSP Settings 40-19
Monitoring CRLs 40-20
Configuring CA Certificates for Revocation 40-20
Configuring CRL Retrieval Policy 40-21
Configuring CRL Retrieval Methods 40-21
Configuring OCSP Rules 40-22
Configuring Advanced CRL and OCSP Settings 40-23
Configuring Identity Certificates Authentication 40-24
Adding or Importing an Identity Certificate 40-24
Showing Identity Certificate Details 40-26
Deleting an Identity Certificate 40-27
Exporting an Identity Certificate 40-27

Generating a Certificate Signing Request 40-27
Installing Identity Certificates 40-28
Configuring Code Signer Certificates 40-29
Showing Code Signer Certificate Details 40-30
Deleting a Code Signer Certificate 40-30
Importing a Code Signer Certificate 40-30
Exporting a Code Signer Certificate 40-31
Authenticating Using the Local CA 40-31
Configuring the Local CA Server 40-32
Deleting the Local CA Server 40-34
Managing the User Database 40-35
Adding a Local CA User 40-35
Sending an Initial OTP or Replacing OTPs 40-36
Editing a Local CA User 40-36
Deleting a Local CA User 40-37
Allowing User Enrollment 40-37
Viewing or Regenerating an OTP 40-37
Managing User Certificates 40-37
Monitoring CRLs 40-38
Feature History for Certificate Management 40-39

PART 8 **Configuring Logging, SNMP, and Smart Call Home**

CHAPTER 41 **Configuring Logging** 41-1
Information About Logging 41-1
Logging in Multiple Context Mode 41-2
Analyzing Syslog Messages 41-2
Syslog Message Format 41-3
Severity Levels 41-3
Message Classes and Range of Syslog IDs 41-4
Filtering Syslog Messages 41-4
Sorting in the Log Viewers 41-4
Using Custom Message Lists 41-5
Using Clustering 41-5
Licensing Requirements for Logging 41-5
Prerequisites for Logging 41-6
Guidelines and Limitations 41-6
Configuring Logging 41-7

Enabling Logging 41-7
Configuring an Output Destination 41-8
Sending Syslog Messages to an External Syslog Server 41-9
Configuring FTP Settings 41-10
Configuring Logging Flash Usage 41-10
Configuring Syslog Messaging 41-10
Editing Syslog ID Settings 41-11
Including a Device ID in Non-EMBLEM Formatted Syslog Messages 41-12
Sending Syslog Messages to the Internal Log Buffer 41-12
Saving an Internal Log Buffer to Flash 41-13
Viewing and Copying Logged Entries with the ASDM Java Console 41-13
Sending Syslog Messages to an E-mail Address 41-14
Adding or Editing E-Mail Recipients 41-14
Configuring the Remote SMTP Server 41-15
Viewing Syslog Messages in ASDM 41-15
Applying Message Filters to a Logging Destination 41-15
Applying Logging Filters 41-16
Adding or Editing a Message Class and Severity Filter 41-17
Adding or Editing a Syslog Message ID Filter 41-17
Sending Syslog Messages to the Console Port 41-17
Sending Syslog Messages to a Telnet or SSH Session 41-18
Creating a Custom Event List 41-18
Generating Syslog Messages in EMBLEM Format to a Syslog Server 41-19
Adding or Editing Syslog Server Settings 41-19
Generating Syslog Messages in EMBLEM Format to Other Output Destinations 41-20
Changing the Amount of Internal Flash Memory Available for Logs 41-20
Configuring the Logging Queue 41-21
Sending All Syslog Messages in a Class to a Specified Output Destination 41-21
Enabling Secure Logging 41-21
Including the Device ID in Non-EMBLEM Format Syslog Messages 41-22
Including the Date and Time in Syslog Messages 41-22
Disabling a Syslog Message 41-22
Changing the Severity Level of a Syslog Message 41-23
Limiting the Rate of Syslog Message Generation 41-23
Assigning or Changing Rate Limits for Individual Syslog Messages 41-24
Adding or Editing the Rate Limit for a Syslog Message 41-24
Editing the Rate Limit for a Syslog Severity Level 41-24
Monitoring the Logs 41-25
Filtering Syslog Messages Through the Log Viewers 41-25
Editing Filtering Settings 41-27

Executing Certain Commands Using the Log Viewers 41-27
Feature History for Logging 41-28

CHAPTER 42 **Configuring SNMP** 42-1
Information About SNMP 42-1
Information About SNMP Terminology 42-2
SNMP Version 3 42-2
SNMP Version 3 Overview 42-3
Security Models 42-3
SNMP Groups 42-3
SNMP Users 42-3
SNMP Hosts 42-3
Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software 42-4
Licensing Requirements for SNMP 42-4
Prerequisites for SNMP 42-4
Guidelines and Limitations 42-4
Configuring SNMP 42-6
Enabling SNMP 42-6
Configuring an SNMP Management Station 42-7
Configuring SNMP Traps 42-7
Using SNMP Version 1 or 2c 42-8
Using SNMP Version 3 42-9
Configuring a Group of Users 42-10
Monitoring SNMP 42-11
SNMP Syslog Messaging 42-11
SNMP Monitoring 42-11
Where to Go Next 42-12
Additional References 42-12
RFCs for SNMP Version 3 42-12
MIBs 42-13
Application Services and Third-Party Tools 42-14
Feature History for SNMP 42-15

CHAPTER 43 **Configuring NetFlow Secure Event Logging (NSEL)** 43-1
Information About NSEL 43-1
Using NSEL and Syslog Messages 43-2
Using NSEL in Clustering 43-3

Licensing Requirements for NSEL 43-4
Prerequisites for NSEL 43-4
Guidelines and Limitations 43-4
Configuring NSEL 43-5
Using NetFlow 43-5
Matching NetFlow Events to Configured Collectors 43-6
Monitoring NSEL 43-7
Where to Go Next 43-7
Additional References 43-7
Related Documents 43-8
RFCs 43-8
Feature History for NSEL 43-8

CHAPTER 44 **Configuring Anonymous Reporting and Smart Call Home** 44-1
Information About Anonymous Reporting and Smart Call Home 44-1
Information About Anonymous Reporting 44-1
DNS Requirement 44-2
Anonymous Reporting and Smart Call Home Prompt 44-2
Information About Smart Call Home 44-3
Licensing Requirements for Anonymous Reporting and Smart Call Home 44-3
Prerequisites for Smart Call Home and Anonymous Reporting 44-4
Guidelines and Limitations 44-4
Configuring Anonymous Reporting and Smart Call Home 44-5
Configuring Anonymous Reporting 44-5
Configuring Smart Call Home 44-5
Monitoring Anonymous Reporting and Smart Call Home 44-9
Feature History for Anonymous Reporting and Smart Call Home 44-10

PART 9 **System Administration**

CHAPTER 45 **Configuring Management Access** 45-1
Configuring ASA Access for ASDM, Telnet, or SSH 45-1
Licensing Requirements for ASA Access for ASDM, Telnet, or SSH 45-1
Guidelines and Limitations 45-2
Configuring Management Access 45-3
Using a Telnet Client 45-4
Using an SSH Client 45-5

Configuring CLI Parameters 45-5
Licensing Requirements for CLI Parameters 45-5
Guidelines and Limitations 45-5
Configuring a Login Banner 45-6
Customizing a CLI Prompt 45-7
Changing the Console Timeout 45-8
Configuring ICMP Access 45-8
Information About ICMP Access 45-9
Licensing Requirements for ICMP Access 45-9
Guidelines and Limitations 45-9
Default Settings 45-9
Configuring ICMP Access 45-10
Configuring Management Access Over a VPN Tunnel 45-10
Licensing Requirements for a Management Interface 45-11
Guidelines and Limitations 45-11
Configuring a Management Interface 45-11
Configuring AAA for System Administrators 45-12
Information About AAA for System Administrators 45-12
Information About Management Authentication 45-12
Information About Command Authorization 45-13
Licensing Requirements for AAA for System Administrators 45-15
Prerequisites 45-15
Guidelines and Limitations 45-16
Default Settings 45-17
Configuring Authentication for CLI, ASDM, and enable command Access 45-17
Limiting User CLI and ASDM Access with Management Authorization 45-18
Configuring a Password Policy for Local Database Users 45-20
Configuring the Password Policy 45-21
Changing Your Password 45-22
Configuring Command Authorization 45-23
Configuring Local Command Authorization 45-23
Viewing Local Command Privilege Levels 45-25
Configuring Commands on the TACACS+ Server 45-25
Configuring TACACS+ Command Authorization 45-27
Configuring Management Access Accounting 45-28
Viewing the Currently Logged-In User 45-29
Setting a Management Session Quota 45-30
Recovering from a Lockout 45-31
Monitoring Device Access 45-32

Feature History for Management Access 45-33

CHAPTER 46 **Managing Software and Configurations 46-1**

Upgrading the Software 46-1
Upgrade Path and Migrations 46-1
Viewing Your Current Version 46-3
Downloading the Software from Cisco.com 46-3
Upgrading a Standalone Unit 46-3
Upgrading from Your Local Computer 46-4
Upgrading Using the Cisco.com Wizard 46-5
Upgrading a Failover Pair or ASA Cluster 46-6
Upgrading an Active/Standby Failover Pair 46-6
Upgrading an Active/Active Failover Pair 46-8
Upgrading an ASA Cluster 46-10
Managing Files 46-12
Configuring File Access 46-13
Configuring the FTP Client Mode 46-13
Configuring the ASA as a Secure Copy Server 46-13
Customizing the ASA Secure Copy Client 46-14
Configuring the ASA TFTP Client Path 46-15
Adding Mount Points 46-15
Accessing the File Management Tool 46-17
Transferring Files 46-18
Transferring Files Between Local PC and Flash 46-18
Transferring Files Between Remote Server and Flash 46-18
Configuring the Images and Startup Configuration to Use 46-19
Backing Up and Restoring Configurations or Other Files 46-21
Backing Up Configurations 46-21
Backing Up the Local CA Server 46-24
Restoring Configurations 46-25
Saving the Running Configuration to a TFTP Server 46-28
Scheduling a System Restart 46-28
Downgrading Your Software 46-29
Information About Activation Key Compatibility 46-29
Performing the Downgrade 46-30
Configuring Auto Update 46-31
Information About Auto Update 46-31
Auto Update Client or Server 46-31
Auto Update Benefits 46-31

Auto Update Server Support in Failover Configurations 46-32
Guidelines and Limitations 46-34
Configuring Communication with an Auto Update Server 46-34
Feature History for Software and Configurations 46-36

CHAPTER 47 **Troubleshooting** 47-1
Configuring and Running Captures with the Packet Capture Wizard 47-1
Ingress Traffic Selector 47-3
Egress Traffic Selector 47-4
Buffers 47-4
Summary 47-4
Run Captures 47-4
Save Captures 47-5

PART 10 **Reference**

APPENDIX 48 **Addresses, Protocols, and Ports** 48-1
IPv4 Addresses and Subnet Masks 48-1
Classes 48-1
Private Networks 48-2
Subnet Masks 48-2
Determining the Subnet Mask 48-3
Determining the Address to Use with the Subnet Mask 48-3
IPv6 Addresses 48-5
IPv6 Address Format 48-5
IPv6 Address Types 48-6
Unicast Addresses 48-6
Multicast Address 48-8
Anycast Address 48-9
Required Addresses 48-10
IPv6 Address Prefixes 48-10
Protocols and Applications 48-11
TCP and UDP Ports 48-11
Local Ports and Protocols 48-14
ICMP Types 48-15

INDEX



About This Guide

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

- Document Objectives, page xli
- Related Documentation, page xli
- Conventions, page xlii
- Obtaining Documentation and Submitting a Service Request, page xlii

Document Objectives

The purpose of this guide is to help you configure general operations for 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*.

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

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and 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

Released: December 3, 2012
Updated: March 31, 2014

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 Compatibility*.

This chapter includes the following sections:

- ASDM Requirements, page 1-1
- Hardware and Software Compatibility, page 1-7
- VPN Compatibility, page 1-7
- New Features, page 1-7
- Firewall Functional Overview, page 1-18
- VPN Functional Overview, page 1-23
- Security Context Overview, page 1-23
- ASA Clustering Overview, page 1-24

ASDM Requirements

- ASDM Client Operating System and Browser Requirements, page 1-2
- Java and Browser Compatibility, page 1-3

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): • 8[1] • 7 • Vista • 2008 Server • XP	6 through 10. Version 11 or later is not supported.	1.5 or later	No support	18 or later	6 or later
Apple Macintosh OS X 10.4 and later.	No support	1.5 or later	2 or later	18 or later	6 or later
Red Hat Enterprise Linux 5 (GNOME or KDE): • Desktop • Desktop with Workstation	N/A	1.5 or later	N/A	18 or later	6 or later

1. ASDM Version 7.1(3) and later.

Java and Browser Compatibility

Table 1-2 lists compatibility caveats for Java, ASDM, and browser compatibility.

Table 1-2 Caveats for ASDM Compatibility

Java Version	Conditions	Notes
7 update 51	• ASDM Launcher requires trusted certificate • Java Web Start requires newer ASDM version *or* workaround	To continue using the Launcher, do one of the following: • Install a trusted certificate on the ASA from a known CA. • Install a self-signed certificate and register it with Java. See the ASDM certificate procedure in this document. • Downgrade Java to 7 update 45 or earlier. • Alternatively use Java Web Start. To use Java Web Start, do one of the following: • Upgrade ASDM to Version 7.1(5.100) or later. This ASDM version includes the Permissions attribute in the JAR manifest, which is required as of Java 7 Update 51. • To use ASDM 7.1(5) or earlier, add a security exception in the Java Control Panel for each ASA you want to manage with ASDM. See the "Workaround" section at: http://java.com/en/download/help/java_blocked.xml If you already upgraded Java, and can no longer launch ASDM in order to upgrade it to Version 7.1(5.100) or later, then you can either use the CLI to upgrade ASDM, or you can use the above security exception work-around to launch the older ASDM, after which you can upgrade to a newer version.
7 update 45	ASDM shows a yellow warning about the missing Permissions attribute	Java 7 update 45 shows a warning when an application does not have the Permissions attribute in the JAR manifest. **It is safe to ignore this warning**. To prevent this warning from appearing, upgrade to ASDM 7.1(5.100) or later; this ASDM version includes the Permissions attribute, which will be required as of Java 7 Update 51. **Note** Due to a bug in Java, even if you upgrade to ASDM 7.1(5.100) or later, if you also do not have a trusted certificate installed on the ASA, you continue to see the yellow warning about the missing Permissions attribute. To prevent the warning from appearing, install a trusted certificate (from a known CA); or generate a self-signed certificate on the ASA by choosing **Configuration > Device Management > Certificates > Identity Certificates**. Launch ASDM, and when the certificate warning is shown, check the **Always trust connections to websites** checkbox.

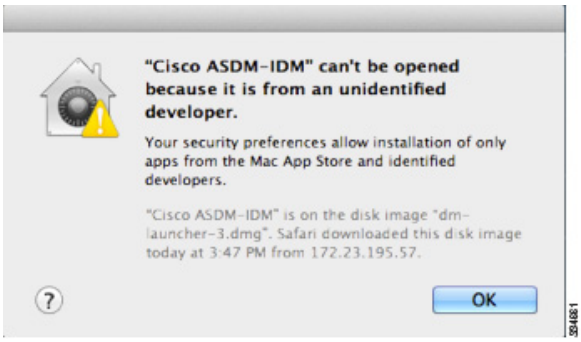
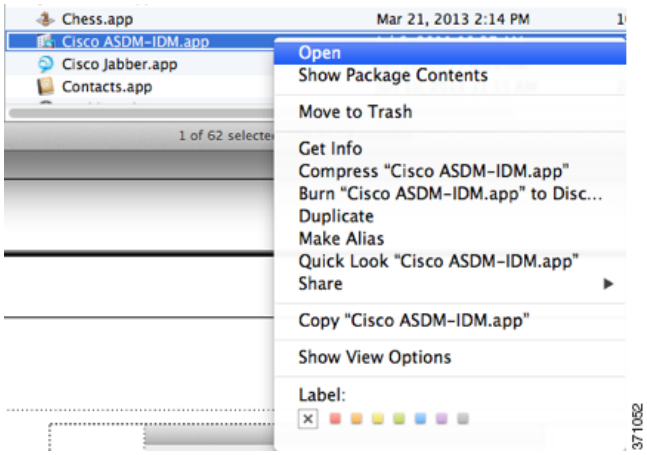
Table 1-2 Caveats for ASDM Compatibility

Java Version	Conditions	Notes
7	Requires strong encryption license (3DES/AES) on ASA	ASDM requires an SSL connection to the ASA. If the ASA has only the base encryption license (DES), and therefore has weak encryption ciphers for the SSL connection, you cannot launch ASDM. You must uninstall Java 7, and install Java 6 (http://www.oracle.com/technetwork/java/javase/downloads/java-archive-downloads-javase6-419409.html). Note that a workaround is required for weak encryption and Java 6 (see below, in this table).
	• ASDM 7.1(3) and earlier • 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. Either upgrade ASDM to 7.1(4) or later, or 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.
6	No usernames longer than 50 characters	Due to a Java bug, ASDM does not support usernames longer than 50 characters when using Java 6. Longer usernames work correctly for Java 7.
	Requires strong encryption license (3DES/AES) on ASA *or* workaround	When you initially connect a browser to the ASA to load the ASDM splash screen, the browser attempts to make an SSL connection to the ASA. If the ASA has only the base encryption license (DES), and therefore has weak encryption ciphers for the SSL connection, you may not be able to access the ASDM splash screen; most current browsers do not support weak encryption ciphers. Therefore, without the strong encryption license (3DES/AES), use one of the following workarounds: • If available, use an already downloaded ASDM launcher or Java Web Start shortcut. The Launcher and Web Start shortcut work 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.

Table 1-2 Caveats for ASDM Compatibility

Java Version	Conditions	Notes
All	• Self-signed certificate or an untrusted certificate • IPv6 • Firefox and Safari	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.
	• SSL encryption on the ASA must include both RC4-MD5 and RC4-SHA1 *or* disable SSL false start in Chrome. • Chrome	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.
	IE9 for servers	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.
	MacOS	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.

Table 1-2 Caveats for ASDM Compatibility

Java Version	Conditions	Notes
All	MacOS 10.8 and later	You need to allow ASDM to run because it is not signed with an Apple Developer ID. If you do not change your security preferences, you see an error screen.  1. To allow ASDM to run, right-click (or Ctrl-Click) the **Cisco ASDM-IDM Launcher** icon, and choose **Open**.  2. You see a similar error screen; however, you can open ASDM from this screen. Click **Open**. The ASDM-IDM Launcher opens. 

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 Compatibility

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 9.1(5)/ASDM 7.1(6), page 1-7
- New Features in ASA 9.1(4)/ASDM 7.1(5), page 1-8
- New Features in ASA 9.1(3)/ASDM 7.1(4), page 1-10
- New Features in ASA 9.1(2)/ASDM 7.1(3), page 1-12
- New Features in ASA 8.4(6)/ASDM 7.1(2.102), page 1-18
- New Features in ASA 9.0(2)/ASDM 7.1(2), page 1-18
- New Features in ASA 9.1(1)/ASDM 7.1(1), page 1-18

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

New Features in ASA 9.1(5)/ASDM 7.1(6)

Released: March 31, 2014

Table 1-3 lists the new features for ASA Version 9.1(5)/ASDM Version 7.1(6).

Table 1-3 New Features for ASA Version 9.1(5)/ASDM Version 7.1(6)

Feature	Description
Administrative Features	
Secure Copy client	The ASA now supports the Secure Copy (SCP) client to transfer files to and from a SCP server. We modified the following screens: Tools > File Management > File Transfer > Between Remote Server and Flash Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server

Table 1-3 New Features for ASA Version 9.1(5)/ASDM Version 7.1(6) (continued)

Feature	Description
Improved one-time password authentication	Administrators who have sufficient authorization privileges may enter privileged EXEC mode by entering their authentication credentials once. The **auto-enable** option was added to the **aaa authorization exec** command. We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization.
Firewall Features	
Transactional Commit Model on rule engine for access groups	When enabled, a rule update is applied after the rule compilation is completed; without affecting the rule matching performance. We introduced the following screen: Configuration > Device Management > Advanced > Rule Engine.
Monitoring Features	
SNMP hosts, host groups, and user lists	You can now add up to 4000 hosts. The number of supported active polling destinations is 128. You can specify a network object to indicate the individual hosts that you want to add as a host group. You can associate more than one user with one host. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
Remote Access Features	
AnyConnect DTLS Single session Performance Improvement	UDP traffic, such as streaming media, was being affected by a high number of dropped packets when sent over an AnyConnect DTLS connection. For example, this could result in streaming video playing poorly or cease streaming completely. The reason for this was the relatively small size of the flow control queue. We increased the DTLS flow-control queue size and offset this by reducing the admin crypto queue size. For TLS sessions, the priority of the crypto command was increased to high to compensated for this change. For both DTLS and TLS sessions, the session will now persist even if packets are dropped. This will prevent media streams from closing and ensure that the number of dropped packets is comparable with other connection methods. We did not modify any ASDM screens.
Webtype ACL enhancements	We introduced URL normalization. URL normalization is an additional security feature that includes path normalization, case normalization and scheme normalization. URLs specified in an ACE and portal address bar are normalized before comparison; for making decisions on webvpn traffic filtering. We did not modify any ASDM screens.

New Features in ASA 9.1(4)/ASDM 7.1(5)

Released: December 9, 2013

Table 1-4 lists the new features for ASA Version 9.1(4)/ASDM Version 7.1(5).

Table 1-4 New Features for ASA Version 9.1(4)/ASDM Version 7.1(5)

Feature	Description
Remote Access Features	
HTML5 WebSocket proxying	HTML5 WebSockets provide persistent connections between clients and servers. During the establishment of the clientless SSL VPN connection, the handshake appears to the server as an HTTP Upgrade request. The ASA will now proxy this request to the backend and provide a relay after the handshake is complete. Gateway mode is not currently supported. We did not modify any ASDM screens.
Inner IPv6 for IKEv2	IPv6 traffic can now be tunneled through IPsec/IKEv2 tunnels. This makes the ASA to AnyConnect VPN connections fully IPv6 compliant. GRE is used when both IPv4 and IPv6 traffic are being tunneled, and when both the client and headend support GRE. For a single traffic type, or when GRE is not supported by the client or the headend, we use straight IPsec. **Note** This feature requires AnyConnect Client Version 3.1.05 or later. We did not modify any ASDM screens.
Mobile Devices running Citrix Server Mobile have additional connection options	Support for mobile devices connecting to Citrix server through the ASA now includes selection of a tunnel-group, and RSA Securid for authorization. Allowing mobile users to select different tunnel-groups allows the administrator to use different authentication methods. We modified the following screen: Configuration > Remote Access VPN > Clientless SSL VPN Access > VDI Access.
Split-tunneling supports exclude ACLs	Split-tunneling of VPN traffic has been enhanced to support both exclude and include ACLs. Exclude ACLs were previously ignored. **Note** This feature requires AnyConnect Client Version 3.1.03103 or later. We did not modify any ASDM screens.
High Availability and Scalability Features	
ASA 5500-X support for clustering	The ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X now support 2-unit clusters. Clustering for 2 units is enabled by default in the base license; for the ASA 5512-X, you need the Security Plus license. We did not modify any ASDM screens.
Improved VSS and vPC support for health check monitoring	If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, you can now increase stability with health check monitoring. For some switches, such as the Nexus 5000, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable the VSS/vPC health check feature, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster

Table 1-4 New Features for ASA Version 9.1(4)/ASDM Version 7.1(5) (continued)

Feature	Description
Support for cluster members at different geographical locations (inter-site); Individual Interface mode only	You can now place cluster members at different geographical locations when using individual interface mode. See the configuration guide for inter-site guidelines. We did not modify any ASDM screens.
Support for clustering with the Cisco Nexus 5000 and Cisco Catalyst 3750-X	The ASA supports clustering when connected to the Cisco Nexus 5000 and Cisco Catalyst 3750-X. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster
Basic Operation Features	
DHCP rebind function	During the DHCP rebind phase, the client now attempts to rebind to other DHCP servers in the tunnel group list. Prior to this release, the client did not rebind to an alternate server, when the DHCP lease fails to renew. We introduced the following screen: Monitoring > Interfaces > DHCP> DHCP Lease Information.
Troubleshooting Features	
Crashinfo dumps include AK47 framework information	Application Kernel Layer 4 to 7 (AK47) framework-related information is now available in crashinfo dumps. A new option, **ak47**, has been added to the **debug menu** command to help in debugging AK47 framework issues. The framework-related information in the crashinfo dump includes the following: • Creating an AK47 instance. • Destroying an AK47 instance. • Generating an crashinfo with a memory manager frame. • Generating a crashinfo after fiber stack overflow. • Generating a crashinfo after a local variable overflow. • Generating a crashinfo after an exception has occurred.

New Features in ASA 9.1(3)/ASDM 7.1(4)

Released: September 18, 2013

Table 1-5 lists the new features for ASA Version 9.1(3)/ASDM Version 7.1(4).

Table 1-5 New Features for ASA Version 9.1(3)/ASDM Version 7.1(4)

Feature	Description
Module Features	
Support for the ASA CX module in multiple context mode	You can now configure ASA CX service policies per context on the ASA. **Note** Although you can configure per context ASA service policies, the ASA CX module itself (configured in PRSM) is a single context mode device; the context-specific traffic coming from the ASA is checked against the common ASA CX policy. Requires ASA CX 9.2(1) or later. We did not modify any ASDM screens.
ASA 5585-X with SSP-40 and -60 support for the ASA CX SSP-40 and -60	ASA CX SSP-40 and -60 modules can be used with the matching level ASA 5585-X with SSP-40 and -60. Requires ASA CX 9.2(1) or later. We did not modify any screens.
Filtering packets captured on the ASA CX backplane	You can now filter packets that have been captured on the ASA CX backplane using the **match** or **access-list** keyword with the **capture interface asa_dataplane** command. Control traffic specific to the ASA CX module is not affected by the access-list or match filtering; the ASA captures all control traffic. In multiple context mode, configure the packet capture per context. Note that all control traffic in multiple context mode goes only to the system execution space. Because only control traffic cannot be filtered using an access list or match, these options are not available in the system execution space. Requires ASA CX 9.2(1) or later. A new option, Use backplane channel, was added to the Ingress Traffic Selector screen and the Egress Selector screen, in the Packet Capture Wizard to enable filtering of packets that have been captured on the ASA CX backplane.
Monitoring Features	
Smart Call Home	We added a new type of Smart Call Home message to support ASA clustering. A Smart Call Home clustering message is sent for only the following three events: • When a unit joins the cluster • When a unit leaves the cluster • When a cluster unit becomes the cluster master Each message that is sent includes the following information: • The active cluster member count • The output of the **show cluster info** command and the **show cluster history** command on the cluster master We did not modify any ASDM screens. *Also available in 9.0(3).*

Table 1-5 New Features for ASA Version 9.1(3)/ASDM Version 7.1(4) (continued)

Feature	Description
Remote Access Features	
user-storage value command password is now encrypted in **show** commands	The password in the **user-storage value** command is now encrypted when you enter **show running-config**. We modified the following screen: **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies > More Options > Session Settings**. *Also available in 8.4(6).*

New Features in ASA 9.1(2)/ASDM 7.1(3)

Released: May 14, 2013

Table 1-6 lists the new features for ASA Version 9.1(2)/ASDM Version 7.1(3).

Note Features added in 8.4(6) are not included in 9.1(2) unless they are explicitly listed in this table.

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3)

Feature	Description
Certification Features	
FIPS and Common Criteria certifications	The FIPS 140-2 Non-Proprietary Security Policy was updated as part of the Level 2 FIPS 140-2 validation for the Cisco ASA series, which includes the Cisco ASA 5505, ASA 5510, ASA 5520, ASA 5540, ASA 5550, ASA 5580, ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X, ASA 5585-X, and the ASA Services Module. The Common Criteria Evaluation Assurance Level 4 (EAL4) was updated, which provides the basis for a specific Target of Evaluation (TOE) of the Cisco ASA and VPN platform solutions.
Encryption Features	
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	Instead of using the proprietary encryption for the failover key, you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3) (continued)

Feature	Description
Additional ephemeral Diffie-Hellman ciphers for SSL encryption	The ASA now supports the following ephemeral Diffie-Hellman (DHE) SSL cipher suites: • DHE-AES128-SHA1 • DHE-AES256-SHA1 These cipher suites are specified in RFC 3268, *Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS*). When supported by the client, DHE is the preferred cipher because it provides Perfect Forward Secrecy. See the following limitations: • DHE is not supported on SSL 3.0 connections, so make sure to also enable TLS 1.0 for the SSL server. • Some popular applications do not support DHE, so include at least one other SSL encryption method to ensure that a cipher suite common to both the SSL client and server can be used. • Some clients may not support DHE, including AnyConnect 2.5 and 3.0, Cisco Secure Desktop, and Internet Explorer 9.0. We modified the following screen: Configuration > Device Management > Advanced > SSL Settings. *Also available in 8.4(4.1).*
Management Features	
Support for administrator password policy when using the local database	When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters. We introduced the following screen: Configuration > Device Management > Users/AAA > Password Policy. *Also available in 8.4(4.1).*
Support for SSH public key authentication	You can now enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF *Also available in 8.4(4.1); PKF key format support is only in 9.1(2).*
AES-CTR encryption for SSH	The SSH server implementation in the ASA now supports AES-CTR mode encryption.
Improved SSH rekey interval	An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic.

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3) (continued)

Feature	Description
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported. We modified the following screen: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH. *Also available in 8.4(4.1).*
Support for a maximum number of management sessions	You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions. We introduced the following screen: Configuration > Device Management > Management Access > Management Session Quota. *Also available in 8.4(4.1).*
Support for a pre-login banner in ASDM	Administrator can define a message that appears before a user logs into ASDM for management access. This customizable content is called a pre-login banner, and can notify users of special requirements or important information.
The default Telnet password was removed	To improve security for management access to the ASA, the default login password for Telnet was removed; you must manually set the password before you can log in using Telnet. **Note**: The login password is only used for Telnet if you do not configure Telnet user authentication. Formerly, when you cleared the password, the ASA restored the default of "cisco." Now when you clear the password, the password is removed. The login password is also used for Telnet sessions from the switch to the ASASM (see the **session** command). For initial ASASM access, you must use the **service-module session** command, until you set a login password. We did not modify any ASDM screens. *Also available in 9.0(2).*
Platform Features	
Support for Power-On Self-Test (POST)	The ASA runs its power-on self-test at boot time even if it is not running in FIPS 140-2-compliant mode. Additional tests have been added to the POST to address the changes in the AES-GCM/GMAC algorithms, ECDSA algorithms, PRNG, and Deterministic Random Bit Generator Validation System (DRBGVS).
Improved pseudo-random number generation (PRNG)	The X9.31 implementation has been upgraded to use AES-256 encryption instead of 3DES encryption to comply with the Network Device Protection Profile (NDPP) in single-core ASAs.
Support for image verification	Support for SHA-512 image integrity checking was added. We did not modify any ASDM screens. *Also available in 8.4(4.1).*
Support for private VLANs on the ASA Services Module	You can use private VLANs with the ASASM. Assign the primary VLAN to the ASASM; the ASASM automatically handles secondary VLAN traffic. There is no configuration required on the ASASM for this feature; see the switch configuration guide for more information.

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3) (continued)

Feature	Description
CPU profile enhancements	The **cpu profile activate** command now supports the following: • Delayed start of the profiler until triggered (global or specific thread CPU%) • Sampling of a single thread We did not modify any ASDM screens. *Also available in 8.4(6).*
DHCP Features	
DHCP relay servers per interface (IPv4 only)	You can now configure DHCP relay servers per-interface, so requests that enter a given interface are relayed only to servers specified for that interface. IPv6 is not supported for per-interface DHCP relay. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP trusted interfaces	You can now configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
Module Features	
ASA 5585-X support for network modules	The ASA 5585-X now supports additional interfaces on network modules in slot 1. You can install one or two of the following optional network modules: • ASA 4-port 10G Network Module • ASA 8-port 10G Network Module • ASA 20-port 1G Network Module *Also available in 8.4(4.1).*
ASA 5585-X DC power supply support	Support was added for the ASA 5585-X DC power supply. *Also available in 8.4(5).*
Support for ASA CX monitor-only mode for demonstration purposes	For demonstration purposes only, you can enable monitor-only mode for the service policy, which forwards a copy of traffic to the ASA CX module, while the original traffic remains unaffected. Another option for demonstration purposes is to configure a traffic-forwarding interface instead of a service policy in monitor-only mode. The traffic-forwarding interface sends all traffic directly to the ASA CX module, bypassing the ASA. We modified the following screen: Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA CX Inspection. The traffic-forwarding feature is supported by CLI only.
Support for the ASA CX module and NAT 64	You can now use NAT 64 in conjunction with the ASA CX module. We did not modify any ASDM screens.

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3) (continued)

Feature	Description
NetFlow Features	
Support for NetFlow flow-update events and an expanded set of NetFlow templates	In addition to adding the flow-update events, there are now NetFlow templates that allow you to track flows that experience a change to their IP version with NAT, as well as IPv6 flows that remain IPv6 after NAT. Two new fields were added for IPv6 translation support. Several NetFlow field IDs were changed to their IPFIX equivalents. For more information, see the *Cisco ASA Implementation Note for NetFlow Collectors*.
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 screen: Configuration > Device Management > Management Access > EtherType Rules. *Also available in 8.4(5).*
Decreased the half-closed timeout minimum value to 30 seconds	The half-closed timeout minimum value for both the global timeout and connection timeout was lowered from 5 minutes to 30 seconds to provide better DoS protection. We modified the following screens: Configuration > Firewall > Service Policy Rules > Connection Settings Configuration > Firewall > Advanced > Global Timeouts.
Remote Access Features	
IKE security and performance improvements	The number of IPsec-IKE security associations (SAs) can be limited for IKE v1 now, as well as IKE v2. We modified the following screen: Configuration > Site-to-Site VPN > Advanced > IKE Parameters.
	The IKE v2 Nonce size has been increased to 64 bytes. There are no ASDM screen or CLI changes.
	For IKE v2 on Site-to-Site, a new algorithm ensures that the encryption algorithm used by child IPsec SAs is not higher strength than the parent IKE. Higher strength algorithms will be downgraded to the IKE level. This new algorithm is enabled by default. We recommend that you do not disable this feature. We did not modify any ASDM screens.
	For Site-to-Site, IPsec data-based rekeying can be disabled. We modified the following screen: Configuration > Site-to-Site > IKE Parameters.
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. *Also available in 8.4(5).*

Table 1-6 New Features for ASA Version 9.1(2)/ASDM Version 7.1(3) (continued)

Feature	Description
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. *Also available in 9.0(2).*
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. *Also available in 9.0(2).*
Dynamic Access Policies: Windows 8 Support	ASDM was updated to enable selection of Windows 8 in the DAP Operating System attribute. *Also available in 9.0(2).*
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. We did not modify any ASDM screens. *Also available in 8.4(5).*
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 *Also available in 8.4(5).*

New Features in ASA 8.4(6)/ASDM 7.1(2.102)

/ASDM Version 7.1(2.102)

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

/ASDM Version 7.1(2)

New Features in ASA 9.1(1)/ASDM 7.1(1)

Released: December 3, 2012

Table 1-7 lists the new features for ASA Version 9.1(1)/ASDM Version 7.1(1).

Note Features added in 8.4(4.x), 8.4(5), 8.4(6), and 9.0(2) are not included in 9.1(1) unless they were listed in the 9.0(1) feature table.

Table 1-7 New Features for ASA Version 9.1(1)/ASDM Version 7.1(1)

Feature	Description
Module Features	
Support for the ASA CX SSP for the ASA 5512-X through ASA 5555-X	We introduced support for the ASA CX SSP software module for the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X. The ASA CX software module requires a Cisco solid state drive (SSD) on the ASA. For more information about the SSD, see the ASA 5500-X hardware guide. We did not modify any screens.

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 1-19
- Firewall Mode Overview, page 1-21
- Stateful Inspection Overview, page 1-22

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 1-19
- Applying NAT, page 1-19
- Protecting from IP Fragments, page 1-19
- Using AAA for Through Traffic, page 1-20
- Applying HTTP, HTTPS, or FTP Filtering, page 1-20
- Applying Application Inspection, page 1-20
- Sending Traffic to a Module, page 1-20
- Applying QoS Policies, page 1-20
- Applying Connection Limits and TCP Normalization, page 1-20
- Enabling Threat Detection, page 1-20
- Enabling the Botnet Traffic Filter, page 1-21
- Configuring Cisco Unified Communications, page 1-21

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 consider using Cisco Web Security Appliance, a dedicated web proxy that provides advanced web filtering as well as advanced threat defense, advanced malware protection, application visibility and control, insightful reporting, and secure mobility. For additional information, see http://www.cisco.com/go/websecurity.

You may also the ASA in conjunction with a third-party server running a filtering product such as Websense Enterprise.

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 a Module

If your model supports an add-on module, then you can send traffic to the module for inspection. For more information, see the documentation for your module.

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 in the firewall configuration guide.

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 2-7
- Verifying the Module Installation, page 2-8
- Assigning VLANs to the ASA Services Module, page 2-9
- Using the MSFC as a Directly Connected Router (SVIs), page 2-12
- Configuring the Switch for ASA Failover, page 2-13
- Resetting the ASA Services Module, page 2-14
- Monitoring the ASA Services Module, page 2-14
- Feature History for the Switch for Use with the ASA Services Module, page 2-17

Information About the Switch

- How the ASA Services Module Works with the Switch, page 2-1
- Supported Switch Hardware and Software, page 2-3
- Backplane Connection, page 2-4
- ASA and IOS Feature Interaction, page 2-4

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.

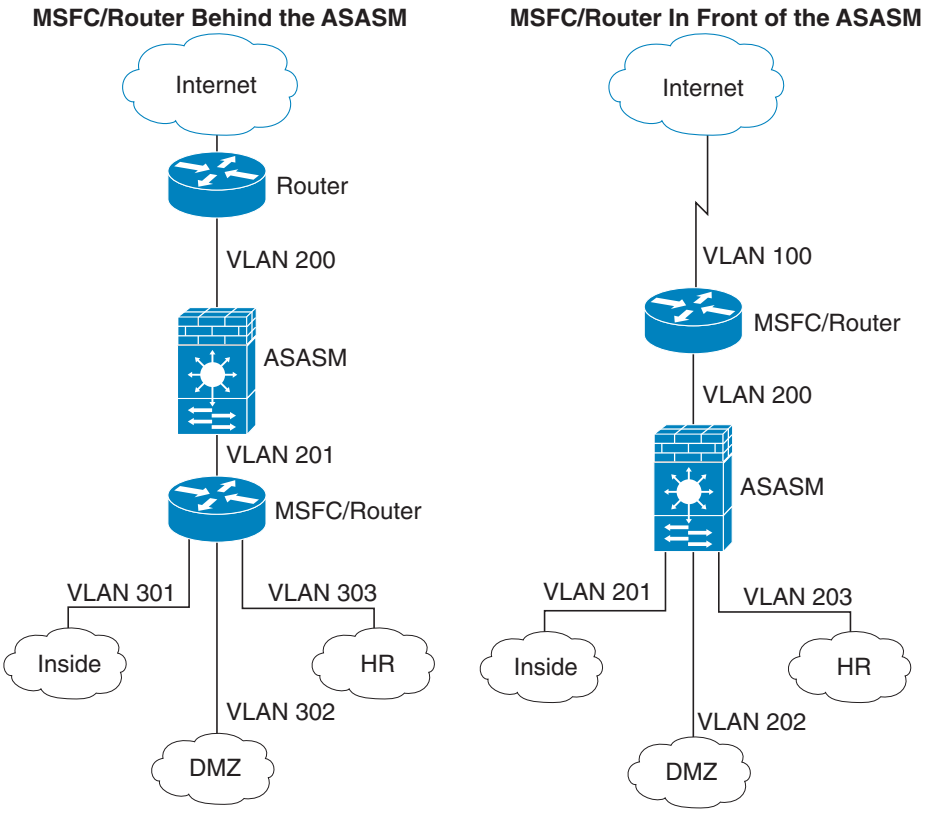
The switch includes a switching processor (the supervisor) and a router (the MSFC). Although you need the MSFC as part of your system, you do not have to use it. If you choose to do so, you can assign one or more VLAN interfaces to the MSFC. You can alternatively use external routers instead of the MSFC.

In single context mode, you can place the router in front of the firewall or behind the firewall (see Figure 2-1).

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 2-1 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 2-1 because you assigned VLAN 200 to the outside interface of the ASASM.

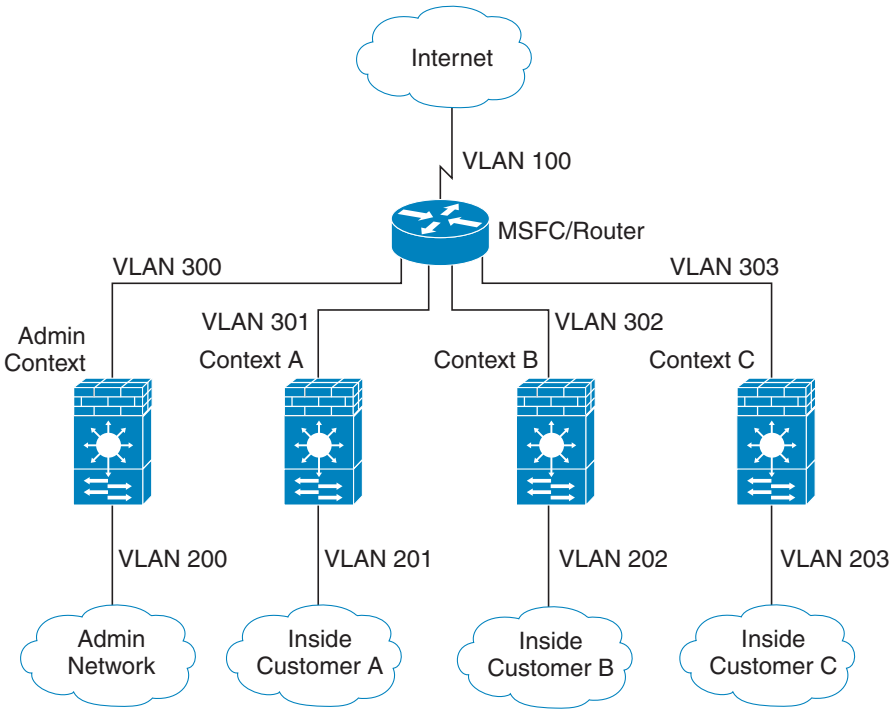
In the left-hand example, the MSFC or router routes between VLANs 201, 301, 302, and 303, and no inside traffic goes through the ASASM unless it is destined for the Internet. In the right-hand example, the ASASM processes and protects all traffic between the inside VLANs 201, 202, and 203.

Figure 2-1 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 2-2).

Figure 2-2 MSFC/Router Placement with Multiple Contexts



Supported Switch Hardware and Software

You can install the ASASM in the Catalyst 6500 series and Cisco 7600 series switches. The switch includes a switch (the supervisor engine) as well as a router (the MSFC).

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

Backplane Connection

The connection between the ASASM and the switch is a single 20-GB interface.

ASA and IOS Feature Interaction

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.

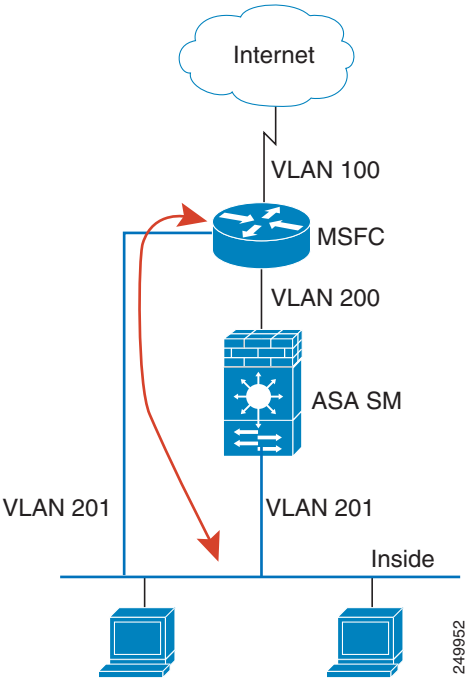
Information About SVIs

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).

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.

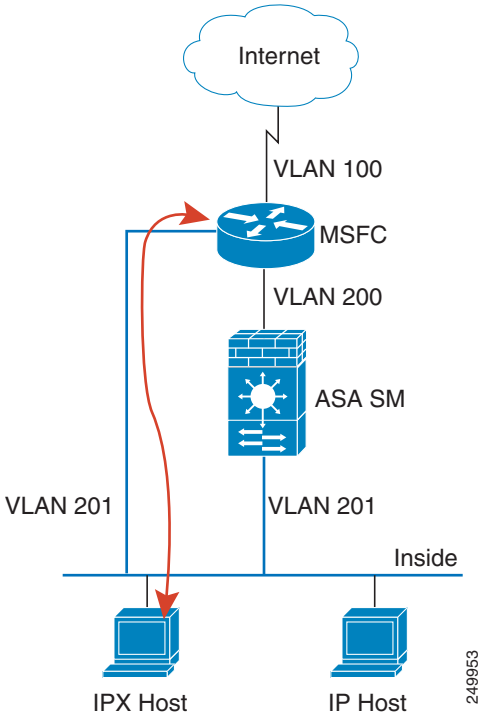
For example, with multiple SVIs, you could accidentally allow traffic to pass around the ASASM by assigning both the inside and outside VLANs to the MSFC. (See Figure 2-3.)

Figure 2-3 Multiple SVI Misconfiguration



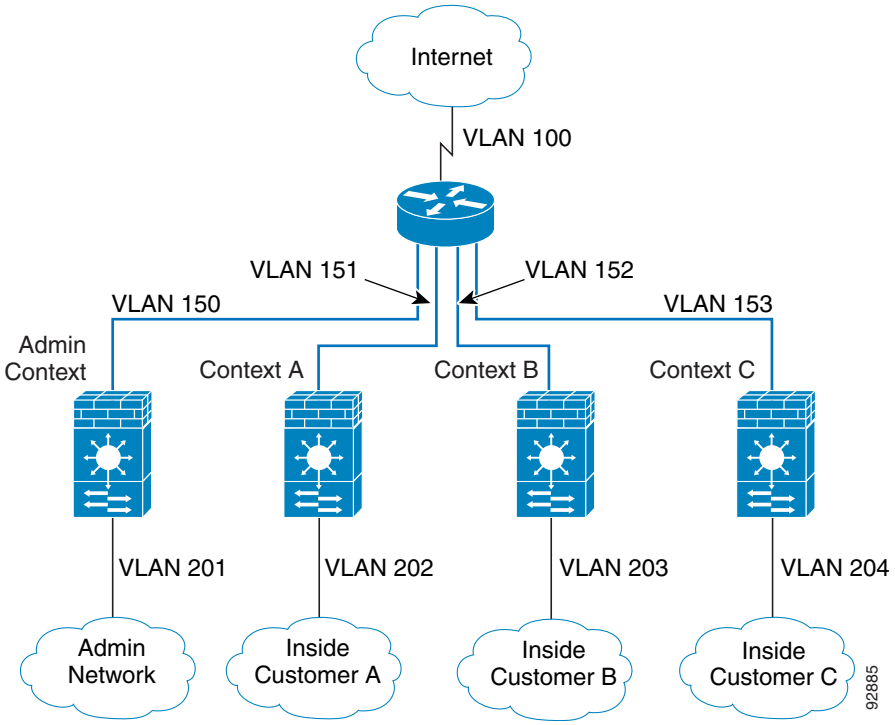
You might need to bypass the ASASM in some network scenarios. Figure 2-4 shows an IPX host on the same Ethernet segment as IP hosts. Because the ASASM in routed firewall mode only handles IP traffic and drops other protocol traffic like IPX (transparent firewall mode can optionally allow non-IP traffic), 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 2-4 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 2-5). 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 2-5 Multiple SVIs in Multiple Context Mode



Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

VLAN Guidelines and Limitations

- Use VLAN IDs 2 to 1001.
- You can use private VLANs with the ASASM. Assign the primary VLAN to the ASASM; the ASASM automatically handles secondary VLAN traffic. There is no configuration required on the ASASM for this feature; see the switch configuration guide for more information. See also the example in "Assigning VLANs to the ASA Services Module" section on page 2-9.
- 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 2-7.

Detailed Steps

	Command	Purpose
Step 1	**firewall vlan-group** *firewall_group vlan_range* **Example:** `Router(config)# firewall vlan-group 1 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 1001) 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 1`	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 10 55-57
Router(config)# firewall vlan-group 11 70-85
Router(config)# firewall vlan-group 12 100
Router(config)# firewall module 5 vlan-group 10,12
Router(config)# firewall module 8 vlan-group 11,12
```

The following example shows how to configure private VLANs on the switch by assigning the primary VLAN to the ASASM:

Step 1 Add the primary VLAN 200 to a firewall VLAN group, and assign the group to the ASASM:

```
Router(config)# firewall vlan-group 10 200
Router(config)# firewall module 5 vlan-group 10
```

Step 2 Designate VLAN 200 as the primary VLAN:

```
Router(config)# vlan 200
Router(config-vlan)# private-vlan primary
```

Step 3 Designate only one secondary isolated VLAN. Designate one or more secondary community VLANs.

```
Router(config)# vlan 501
Router(config-vlan)# private-vlan isolated
Router(config)# vlan 502
Router(config-vlan)# private-vlan community
Router(config)# vlan 503
Router(config-vlan)# private-vlan community
```

Step 4 Associate the secondary VLANs to the primary VLAN:

```
Router(config)# vlan 200
Router(config-vlan)# private-vlan association 501-503
```

Step 5 Classify the port mode. The mode of interface f1/0/1 is host. The mode of interface f1/0/2 is promiscuous.

```
Router(config)# interface f1/0/1
Router(config-ifc)# switchport mode private-vlan host
Router(config)# interface f1/0/2
Router(config-ifc)# switchport mode private-vlan promiscuous
```

Step 6 Assign VLAN membership to the host port. Interface f1/0/1 is a member of primary VLAN 200 and secondary isolated VLAN 501.

```
Router(config)# interface f1/0/1
Router(config-ifc)# switchport private-vlan host-association 200 501
```

Step 7 Assign VLAN membership to the promiscuous interface. Interface f1/0/2 is a member of primary VLAN 200. Secondary VLANs 501-503 are mapped to the primary VLAN.

```
Router(config)# interface f1/0/2
Router(config-ifc)# switchport private-vlan mapping 200 501-503
```

Step 8 If inter-VLAN routing is desired, configure a primary SVI and then map the secondary VLANs to the primary.

```
Router(config)# interface vlan 200
Router(config-ifc)# private-vlan mapping 501-503
```

Using the MSFC as a Directly Connected Router (SVIs)

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). See the "Information About SVIs" section on page 2-5.

Restrictions

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.

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.
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 2-13
- Adding a Trunk Between a Primary Switch and Secondary Switch, page 2-13
- Ensuring Compatibility with Transparent Firewall Mode, page 2-13
- Enabling Autostate Messaging for Rapid Link Failure Detection, page 2-13

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 2-13.

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.

Command	Purpose
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.
Support for private VLANs	9.1(2)	You can use private VLANs with the ASASM. Assign the primary VLAN to the ASASM; the ASASM automatically handles secondary VLAN traffic. There is no configuration required on the ASASM for this feature; see the switch configuration guide for more information.



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 3-7
- Configuring ASDM Access for the ASA Services Module, page 3-12
- Starting ASDM, page 3-14
- Installing an Identity Certificate for ASDM, page 3-18
- Factory Default Configurations, page 3-19
- Getting Started with the Configuration, page 3-25
- Using the Command Line Interface Tool in ASDM, page 3-26
- Applying Configuration Changes to Connections, page 3-28

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 3-7 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 45, "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 8, "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:

```
ciscoasa>
```

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:

```
ciscoasa> 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 16-1 to change the enable password.

The prompt changes to:

```
ciscoasa#
```

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

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

```
ciscoasa# configure terminal
```

The prompt changes to the following:

```
ciscoasa(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 3-5
- 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 8, "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 45-1.

This section includes the following topics:

- Information About Connection Methods, page 3-3
- Logging In, page 3-4

Information About Connection Methods

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

- 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.

  – No initial password configuration is required.

  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 3-5 for more information.

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

  Note You cannot connect using this method for a new ASASM; this method requires you to configure a Telnet login password on the ASASM (there is no default password). After you set a password using the **passwd** command, you can use this method.

  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.
- You must first set a Telnet login password; there is no default password.

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:	

	(Available for initial access.)
	service-module session [**switch** {**1** | **2**}] **slot** *number*
	
	Example:
	`Router# service-module session slot 3`
	`ciscoasa>`
		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.

	(Available after you configure a login password.)
	session [**switch** {**1** |**2**}] **slot** *number* **processor 1**
	You are prompted for the login password:
	`ciscoasa passwd:`
	
	Example:
	`Router# session slot 3 processor 1`
	`ciscoasa passwd: cisco`
	`ciscoasa>`
		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. Set the password using the **passwd** command. 9.1(1): The default password is "cisco." 9.1(2) and later: There is no default password.
		You access user EXEC mode.

	Command	Purpose

Step 2	**enable**
	
	Example:
	`ciscoasa> enable`
	`Password:`
	`ciscoasa#`
		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 16-1.
		
		To exit privileged EXEC mode, enter the **disable**, **exit**, or **quit** command.

Step 3	**configure terminal**
	
	Example:
	`ciscoasa# configure terminal`
	`ciscoasa(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 3-6

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 3-6.

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. Use the **terminal escape-character** *ascii_number* command (to change for this session) or the **default escape-character** *ascii_number* command (to change permanently). For example, to change the sequence for the current session to **Ctrl-w**, **x**, enter **terminal escape-character 23**.

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:

- Using the Pre-login Banner in ASDM, page 3-7
- Accessing ASDM Using the Factory Default Configuration, page 3-7
- Customizing ASDM Access (ASA 5505), page 3-8
- Customizing ASDM Access (ASA 5510 and Higher), page 3-10

Using the Pre-login Banner in ASDM

When you first log in to ASDM, a message may appear even if you have not entered the device IP address, your username, and password. Administrators can create a message that appears before users log into ASDM, which is called a pre-login banner. For example:

```
This device is the property of …. Unauthorized use is not allowed.
```

Click **OK** to close the dialog box that includes the message. Then re-enter the device IP address, your username, and password to open ASDM.

Accessing ASDM Using the Factory Default Configuration

With a factory default configuration (see the "Factory Default Configurations" section on page 3-19), 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 launch ASDM, see the "Starting ASDM" section on page 3-14.

Note To change to multiple context mode, see the "Enabling or Disabling Multiple Context Mode" section on page 8-16. After changing to multiple context mode, you can access ASDM from the admin context using the network settings above.

Customizing ASDM Access (ASA 5505)

Use this procedure if *one or more* of the following conditions applies:

- You do not have a factory default configuration
- You want to change to transparent firewall mode

See also the sample configurations in the "ASA 5505 Default Configuration" section on page 3-21.

Note For routed mode, for quick and easy ASDM access, we recommend applying the factory default configuration with the option to set your own management IP address (see the "Restoring the Factory Default Configuration" section on page 3-19). Use the procedure in this section only if you have special needs such as setting transparent mode, or if you have other configuration that you need to preserve.

Prerequisites

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

Detailed Steps

	Command	Purpose
Step 1	(Optional) **firewall transparent** **Example:** `ciscoasa(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Setting the Firewall Mode (Single Mode)" section on page 6-9 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode:	

	Routed mode:
	interface vlan *number*
	nameif *name*
	security-level *level*
	ip address *ip_address* [*mask*]
	
	Example:
	`ciscoasa(config)# interface vlan 1`
	`ciscoasa(config-if)# nameif inside`
	`ciscoasa(config-if)# security-level 100`
	`ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0`
		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 *number*
	nameif *name*
	security-level *level*
	
	Example:
	`ciscoasa(config)# interface bvi 1`
	`ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0`
	
	`ciscoasa(config)# interface vlan 1`
	`ciscoasa(config-if)# bridge-group 1`
	`ciscoasa(config-if)# nameif inside`
	`ciscoasa(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 0/1**
	switchport access vlan *number*
	no shutdown
	
	Example:
	`ciscoasa(config)# interface ethernet 0/1`
	`ciscoasa(config-if)# switchport access vlan 1`
	`ciscoasa(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:
	`ciscoasa(config)# dhcpd address 192.168.1.5-192.168.1.254 inside`
	`ciscoasa(config)# dhcpd enable inside`
		Sets the DHCP pool for the management network. Make sure you do not include the VLAN interface 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:** `ciscoasa(config)# http server enable`	Enables the HTTP server for ASDM.
Step 6	**http** *ip_address mask interface_name* **Example:** `ciscoasa(config)# http 192.168.1.0 255.255.255.0 inside`	Allows the management host(s) to access ASDM.
Step 7	**write memory** **Example:** `ciscoasa(config)# write memory`	Saves the configuration.
Step 8	To launch ASDM, see the "Starting ASDM" section on page 3-14.	

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
```

Customizing ASDM Access (ASA 5510 and Higher)

Use this procedure if *one or more* of the following conditions applies:

- You do not have a factory default configuration
- You want to change to transparent firewall mode
- You want to change to multiple context mode

Note For routed, single mode, for quick and easy ASDM access, we recommend applying the factory default configuration with the option to set your own management IP address (see the "Restoring the Factory Default Configuration" section on page 3-19). Use the procedure in this section only if you have special needs such as setting transparent or multiple context mode, or if you have other configuration that you need to preserve.

Prerequisites

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

Detailed Steps

	Command	Purpose
Step 1	(Optional) **firewall transparent** **Example:** `ciscoasa(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Setting the Firewall Mode (Single Mode)" section on page 6-9 for more information.

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

Step 3	**dhcpd address** *ip_address-ip_address interface_name*
	dhcpd enable *interface_name*
	
	Example:
	`ciscoasa(config)# dhcpd address 192.168.1.2-192.168.1.254 management`
	`ciscoasa(config)# dhcpd enable management`
		Sets the DHCP pool for the management 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:** `ciscoasa(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:** `ciscoasa(config)# http server enable`	Enables the HTTP server for ASDM.
Step 6	**http** *ip_address mask interface_name* **Example:** `ciscoasa(config)# http 192.168.1.0 255.255.255.0 management`	Allows the management host(s) to access ASDM.
Step 7	**write memory** **Example:** `ciscoasa(config)# write memory`	Saves the configuration.

	Command	Purpose
Step 8	(Optional) **mode multiple** **Example:** `ciscoasa(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 8, "Configuring Multiple Context Mode," for more information.
Step 9	To launch ASDM, see the "Starting ASDM" section on page 3-14.	

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 2-9.
- 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:** `ciscoasa(config)# firewall transparent`	Enables transparent firewall mode. This command clears your configuration. See the "Setting the Firewall Mode (Single Mode)" section on page 6-9 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode:	

	Command	Purpose

	Routed mode:
	interface vlan *number*
	ip address *ip_address* [*mask*]
	nameif *name*
	security-level *level*
	
	Example:
	`ciscoasa(config)# interface vlan 1`
	`ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0`
	`ciscoasa(config-if)# nameif inside`
	`ciscoasa(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.

	Transparent mode:
	interface bvi *number*
	ip address *ip_address* [*mask*]
	
	interface vlan *number*
	bridge-group *bvi_number*
	nameif *name*
	security-level *level*
	
	Example:
	`ciscoasa(config)# interface bvi 1`
	`ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0`
	
	`ciscoasa(config)# interface vlan 1`
	`ciscoasa(config-if)# bridge-group 1`
	`ciscoasa(config-if)# nameif inside`
	`ciscoasa(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:
	`ciscoasa(config)# dhcpd address 192.168.1.2-192.168.1.254 inside`
	`ciscoasa(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:** `ciscoasa(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:** `ciscoasa(config)# http server enable`	Enables the HTTP server for ASDM.

	Command	Purpose
Step 6	**http** *ip_address mask interface_name* **Example:** `ciscoasa(config)# http 192.168.1.0 255.255.255.0 management`	Allows the management host to access ASDM.
Step 7	**write memory** **Example:** `ciscoasa(config)# write memory`	Saves the configuration.
Step 8	(Optional) **mode multiple** **Example:** `ciscoasa(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 8, "Configuring Multiple Context Mode," for more information.
Step 9	To launch ASDM, see the "Starting ASDM" section on page 3-14.	

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 3-15
- Starting ASDM from the ASDM-IDM Launcher, page 3-16
- Starting ASDM from the Java Web Start Application, page 3-16
- Using ASDM in Demo Mode, page 3-17

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, by default 192.168.1.2. See the "Configuring ASDM Access for Appliances" section on page 3-7 or the "Configuring ASDM Access for the ASA Services Module" section on page 3-12 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 3-16 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 3-16 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 3-15.

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 3-15.

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.

Installing an Identity Certificate for ASDM

When using Java 7 update 51 and later, the ASDM Launcher requires a trusted certificate. An easy approach to fulfill the certificate requirements is to install a self-signed identity certificate. You can use Java Web Start to launch ASDM until you install a certificate.

See the following document to install a self-signed identity certificate on the ASA for use with ASDM, and to register the certificate with Java.

http://www.cisco.com/go/asdm-certificate

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 8, "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 3-19
- ASA 5505 Default Configuration, page 3-21
- ASA 5510 and Higher Default Configuration, page 3-25

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

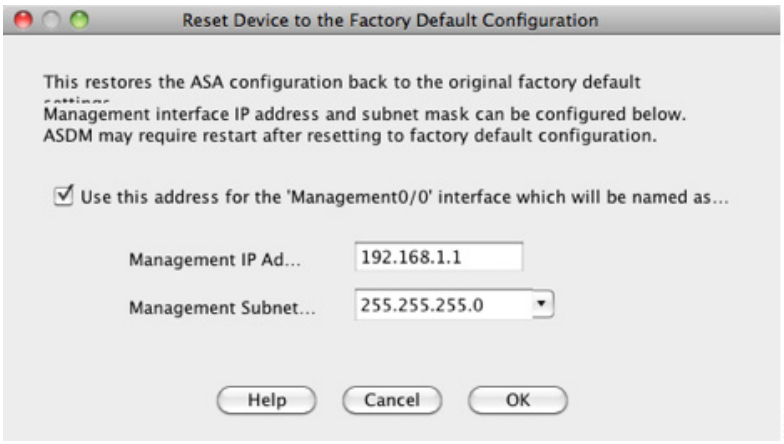
Using the CLI:

	Command	Purpose
Step 1	**configure factory-default** [*ip_address* [*mask*]] **Example:** `ciscoasa(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.

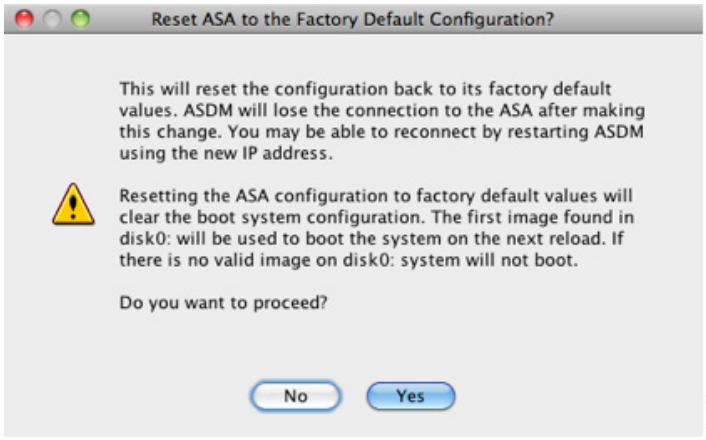


Step 2 (Optional) Enter the Management IP address of the management interface, instead of using the default address, 192.168.1.1. (For an ASA with a dedicated management interface, the interface is called "Management0/0.")

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 3-25 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:

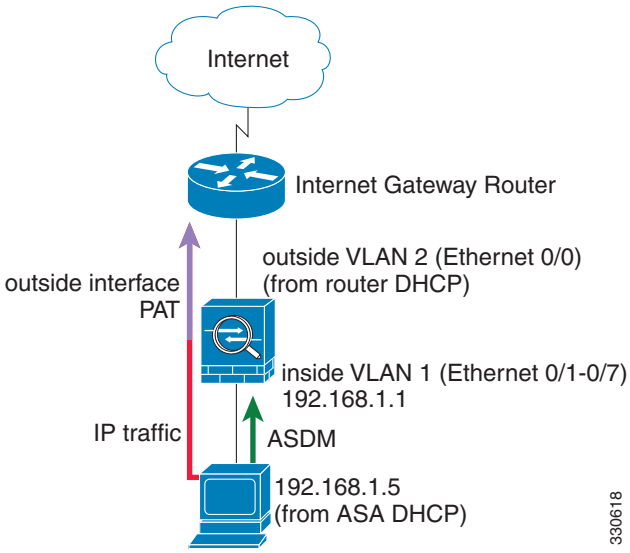
- ASA 5505 Routed Mode Default Configuration, page 3-22
- ASA 5505 Transparent Mode Sample Configuration, page 3-23

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:

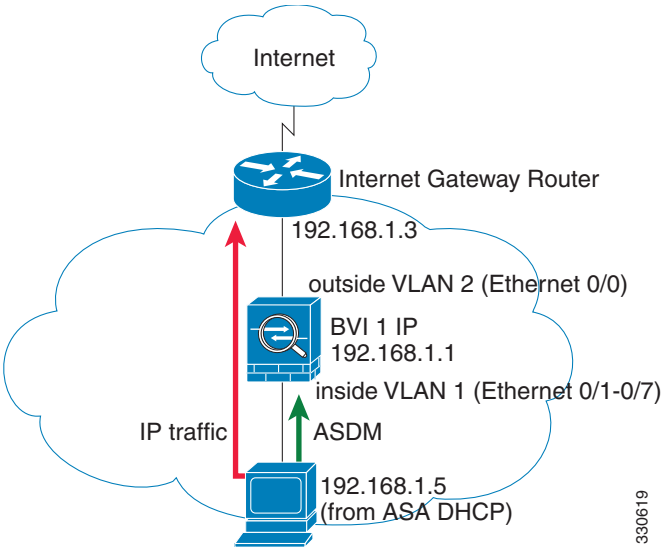
```
policy-map global_policy
   class inspection_default
      inspect icmp
```

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 6-6.

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 3-26
- Handling Command Errors, page 3-27
- Using Interactive Commands, page 3-27
- Avoiding Conflicts with Other Administrators, page 3-27
- Showing Commands Ignored by ASDM on the Device, page 3-27

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 "Authorization" section on page 32-2 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:** `ciscoasa(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:** `ciscoasa(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 4-11
- Status Bar, page 4-11
- Device List, page 4-12
- Common Buttons, page 4-12
- Keyboard Shortcuts, page 4-13
- Find Function, page 4-15
- Enabling Extended Screen Reader Support, page 4-16
- Organizational Folder, page 4-17
- About the Help Window, page 4-17
- 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

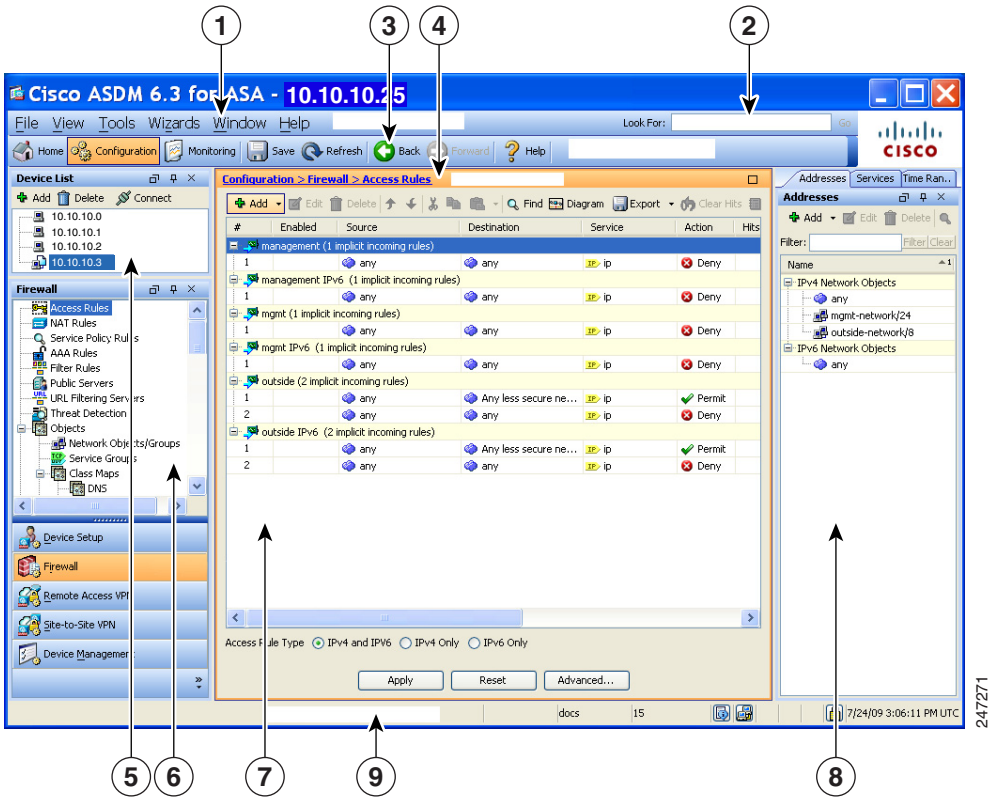
The ASDM user interface is designed to provide easy access to the many features that the ASA supports. The ASDM user interface includes the following elements:

- 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 4-12.
- 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 4-13.

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 3-19 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 46-28 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 4-12 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 4-11 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 3-26 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 3-27 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 in the firewall configuration guide 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 Disabling the Test Configuration section in the firewall configuration guide for more information.
Traceroute	Determines the route that packets will take to their destination. See the Determining Packet Routing with Traceroute section in the firewall configuration guide 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 46-12 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 46-18 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 46-12 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 46-5 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 46-21 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 46-25 for more information.
System Reload	Restarts the ASDM and reload the saved configuration into memory. See the "Scheduling a System Restart" section on page 46-28 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 Capturing Data section in the VPN configuration guide for more information.

Tools Menu Item	Description
Migrate Network Object Group Members	If you migrate to 8.3 or later, the ASA creates named network objects to replace inline IP addresses in some features. In addition to named objects, ASDM automatically creates non-named objects for any IP addresses used in the configuration. These auto-created objects are identified by the *IP address* only, do not have a name, and are not present as named objects in the platform configuration. When the ASA creates named objects as part of the migration, the matching non-named ASDM-only objects are replaced with the named objects. The only exception are non-named objects in a network object group. When the ASA creates named objects for IP addresses that are inside a network object group, ASDM retains the non-named objects as well, creating duplicate objects in ASDM. To merge these objects, choose **Tools > Migrate Network Object Group Members**. See *Cisco ASA 5500 Migration to Version 8.3 and Later* 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 41-13 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 73, "VPN Wizards," in the VPN configuration guide.
SSL VPN Wizard	Enables you to configure an SSL VPN policy on the ASA. For more information, see Chapter 73, "VPN Wizards," in the VPN configuration guide.

Wizards Menu Item	Description
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 "Configuring Active/Standby Failover" section on page 9-25. • Active/Active failover—See the "Configuring Active/Active Failover" section on page 9-32. • VPN cluster load balancing—See the Configuring VPN Cluster Load Balancing with the High Availability and Scalability Wizard section in the VPN configuration guide. • ASA clustering—See the "Adding or Joining an ASA Cluster" section on page 10-50.
Unified Communication Wizard	Enables you to configure unified communication features, such as an IP phone, on the ASA. For more information, see Chapter 15, "Information about the Cisco Unified Communication Wizard," in the firewall configuration guide.
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 47-1.

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.

Help Menu Items	Description
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	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 4-11 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 4-12 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.

Button	Description
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 4-15d
- Using the Find Function in the ACL Manager Pane, page 4-16

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 4-16 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

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.

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

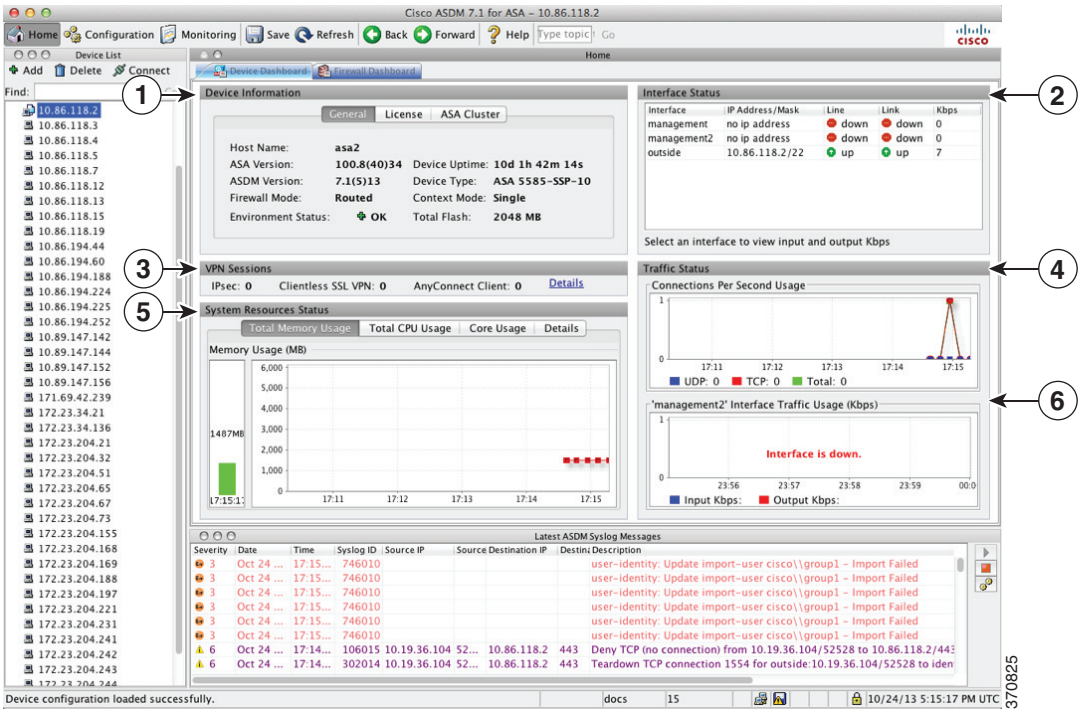
- Content Security Tab, page 4-27
- Intrusion Prevention Tab, page 4-28
- ASA CX Status Tab, page 4-30

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

GUI Element	Description
—	Device List, page 4-12
—	Latest ASDM Syslog Messages Pane, page 4-21

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 4-20
- Cluster Tab, page 4-20

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.

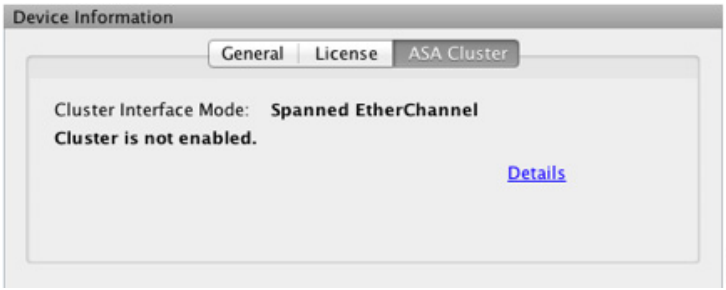
Note If you do not have enough memory to upgrade to the most current release of the ASA, the Memory Insufficient Warning dialog box appears. Follow the directions that appear in this dialog box to continue using the ASA and ASDM in a supported manner. Click **OK** to close this dialog box.

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."

Cluster Tab

This tab shows the cluster interface mode, as well as the cluster status



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.

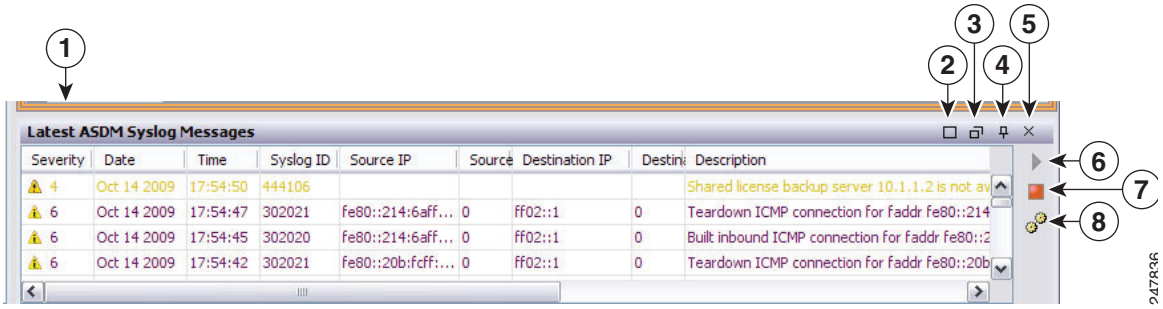
When your configuration contains multiple lowest security level interfaces, and any one of them is named "outside," then that interface is used for the traffic throughput graphs. Otherwise, ASDM picks the first interface from the alphabetical list of lowest security level interfaces.

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.

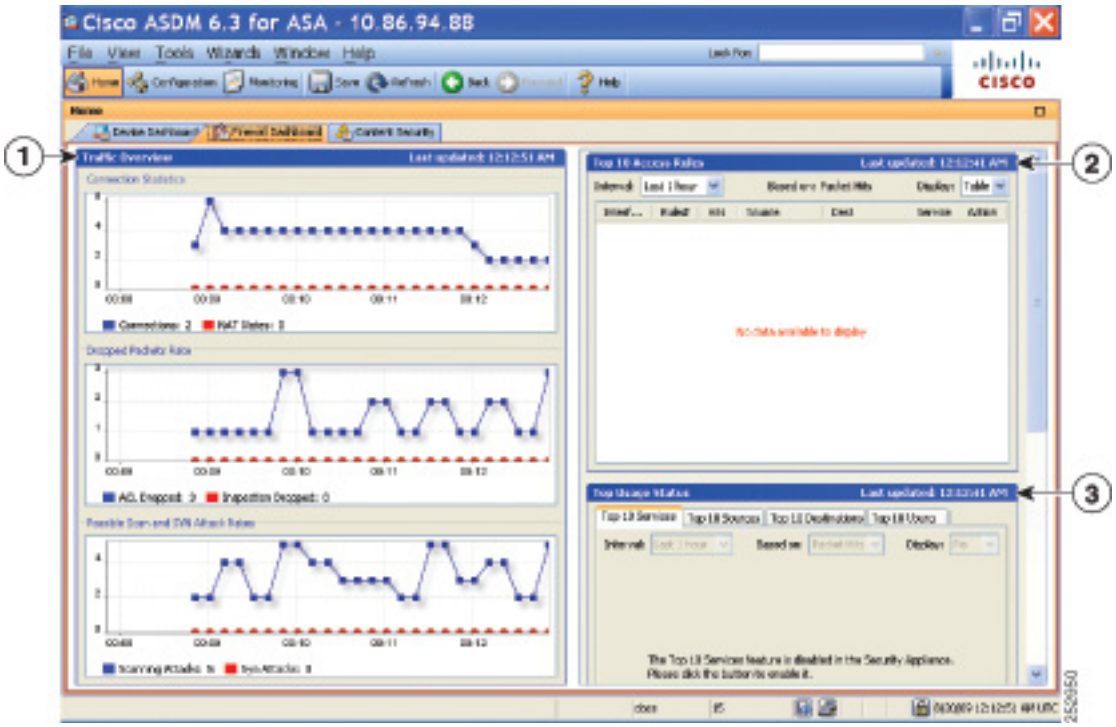
- To clear the current messages, right-click an event and click **Clear Content**.
- To save the current messages to a file on your PC, right-click an event and click **Save Content**.
- To copy the current content, right-click an event and click **Copy**.
- To change the background and foreground colors 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 4-23
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 in the firewall configuration guide), 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 in the firewall configuration guide), 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 in the firewall configuration guide.

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 38-10 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 in the firewall configuration guide. 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.

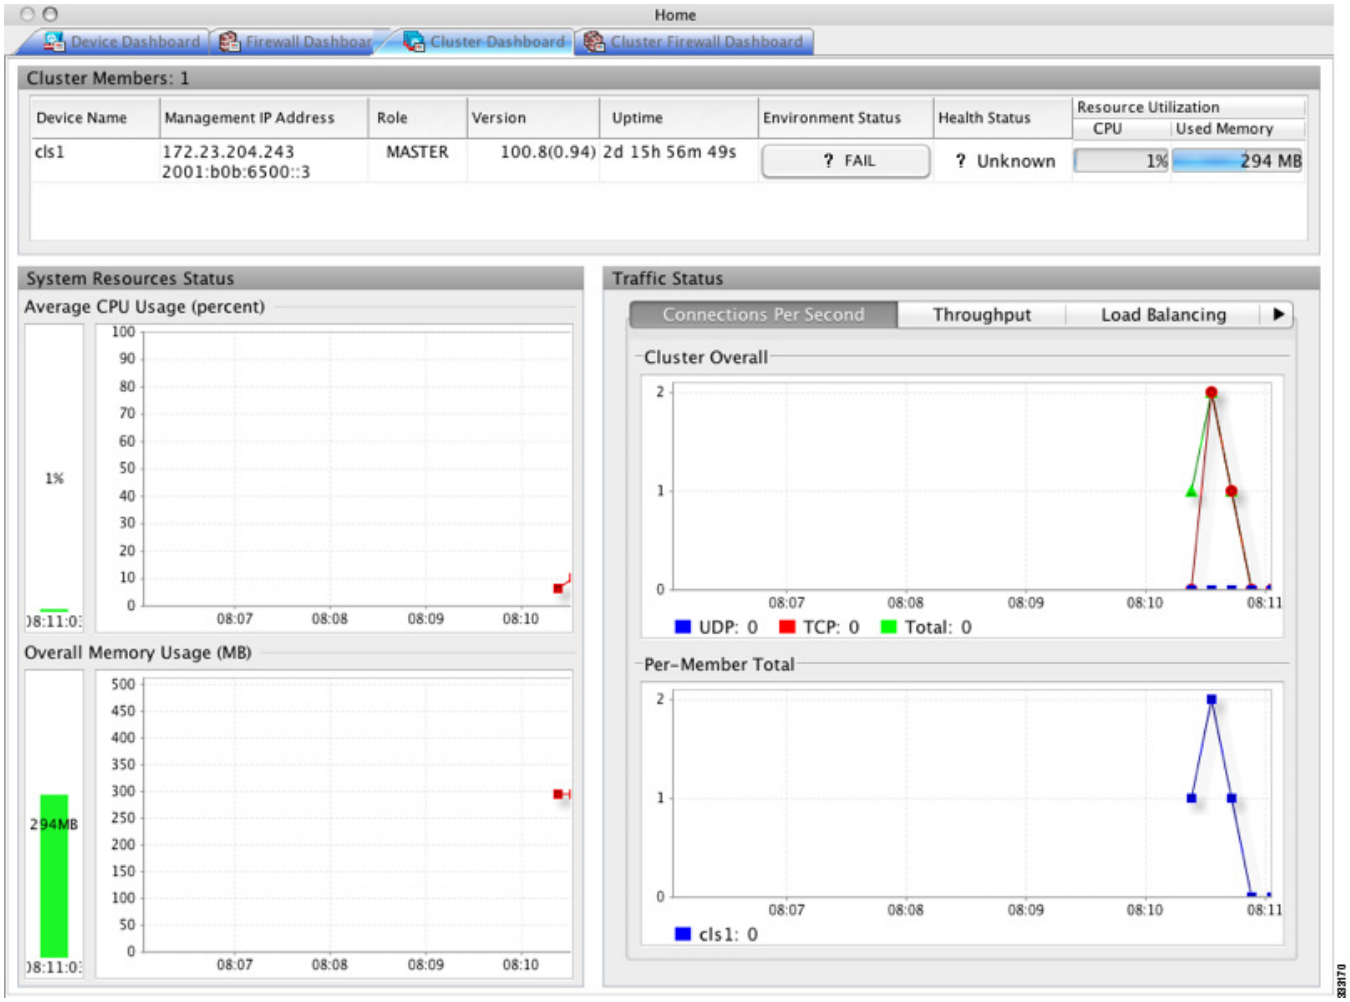
Top Botnet Traffic Filter Hits Pane

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 Chapter 26, "Configuring the Botnet Traffic Filter," in the firewall configuration guide.

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.

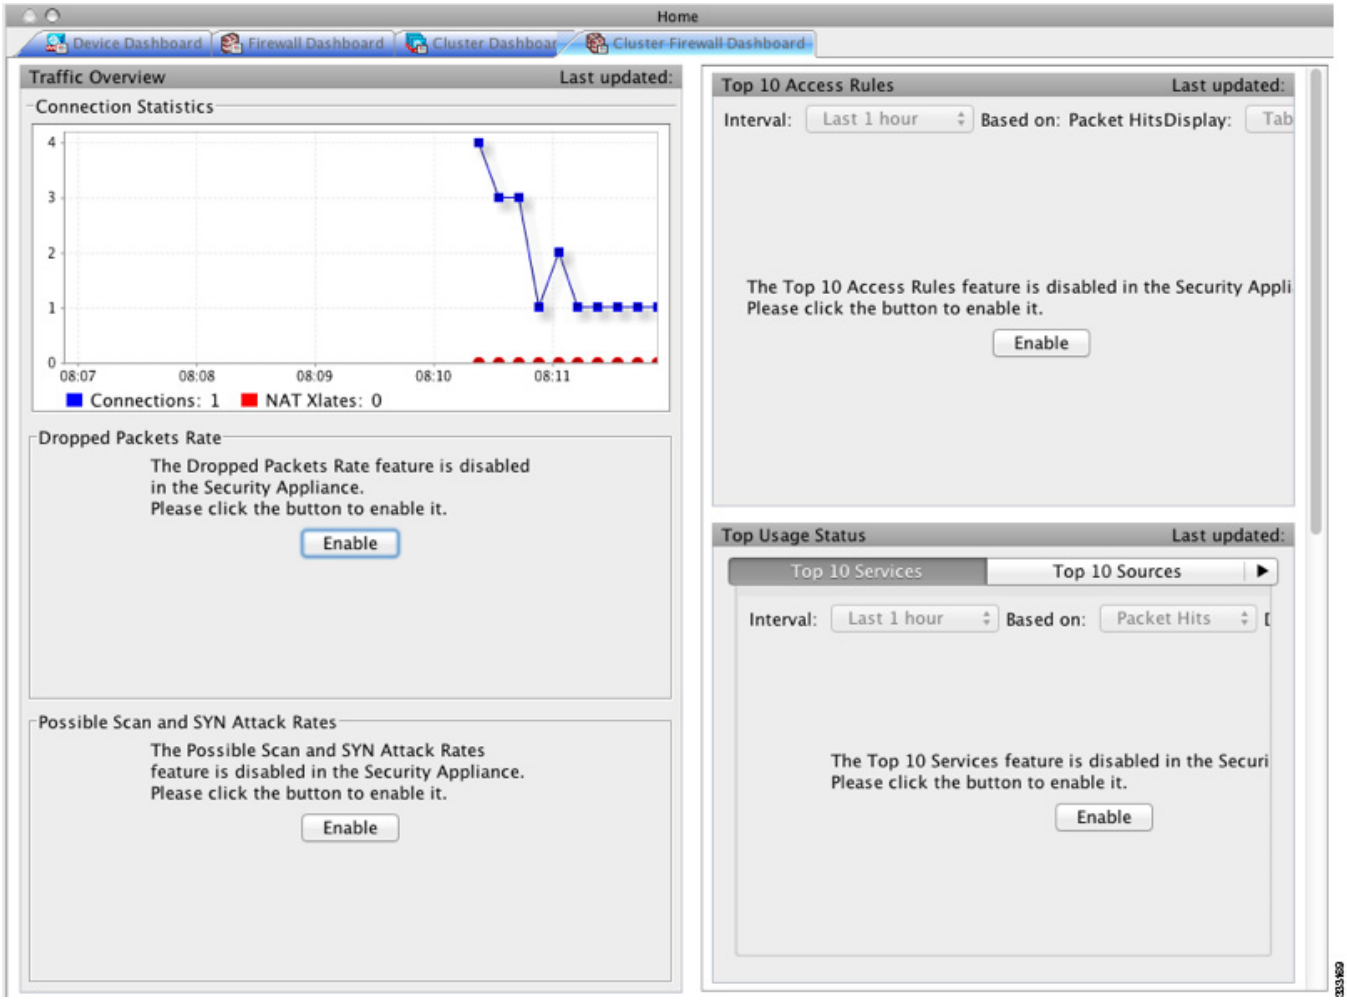
- Control Link Usage tab:

  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

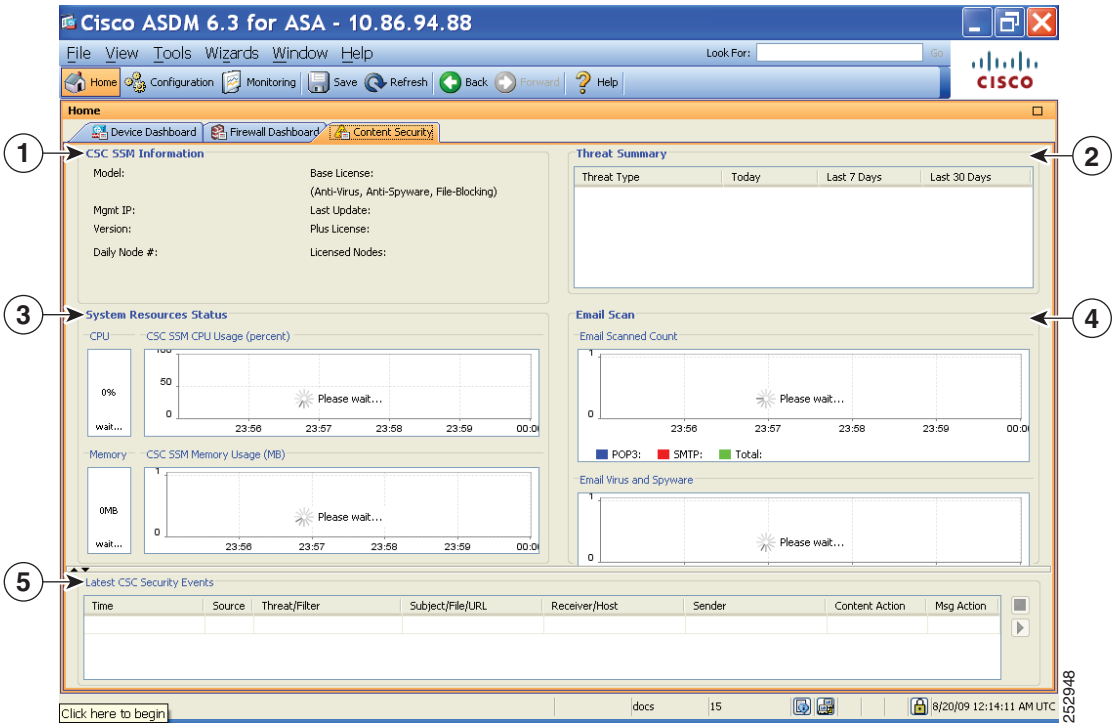
The Content Security tab lets you view important information about the Content Security and Control (CSC) SSM. This pane appears only if CSC software running on the CSC SSM is installed in the ASA.

For an introduction to the CSC SSM, see the Information About the CSC SSM section in the firewall configuration guide.

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.

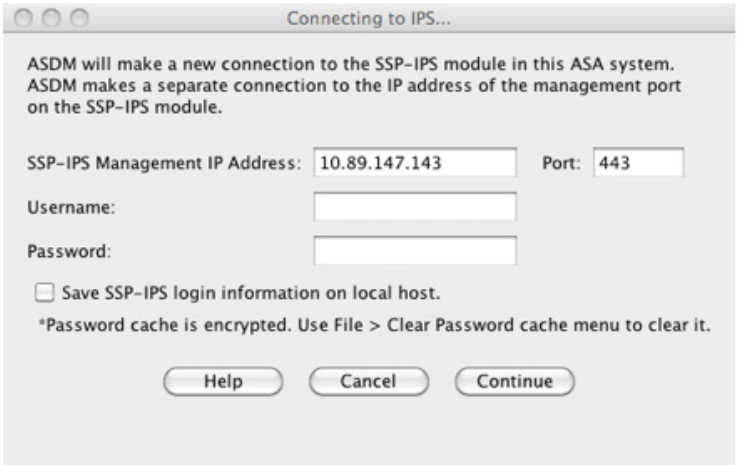
Intrusion Prevention Tab

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**

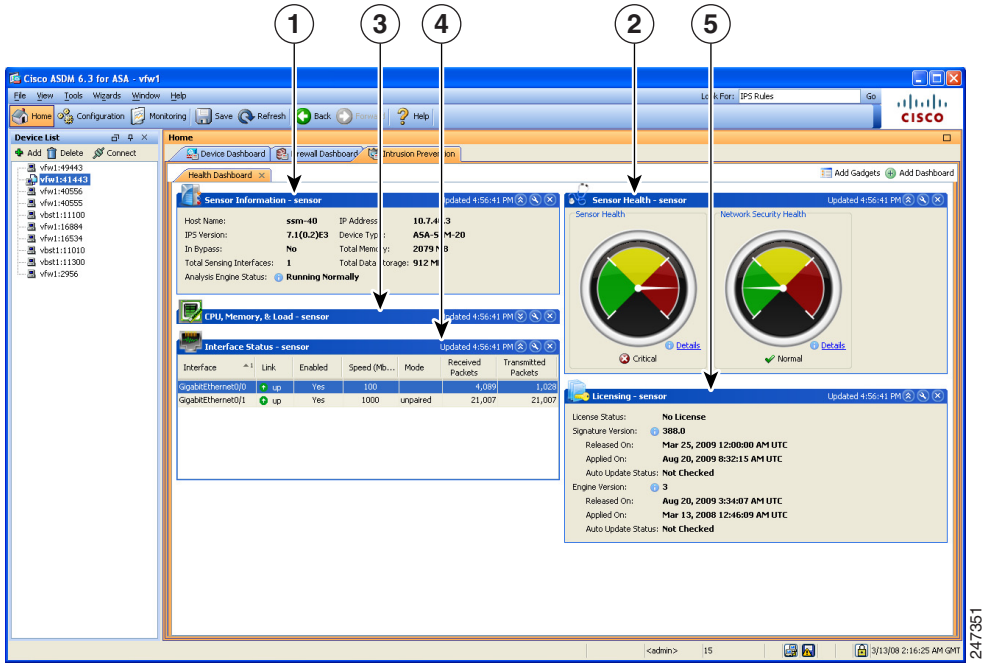
Step 3 To save the login information on your local PC, check the **Save IPS login** information on local host check box.

Step 4 Click **Continue**.

For more information about intrusion prevention, see Chapter 31, "Configuring the ASA IPS Module," in the firewall configuration guide.

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)

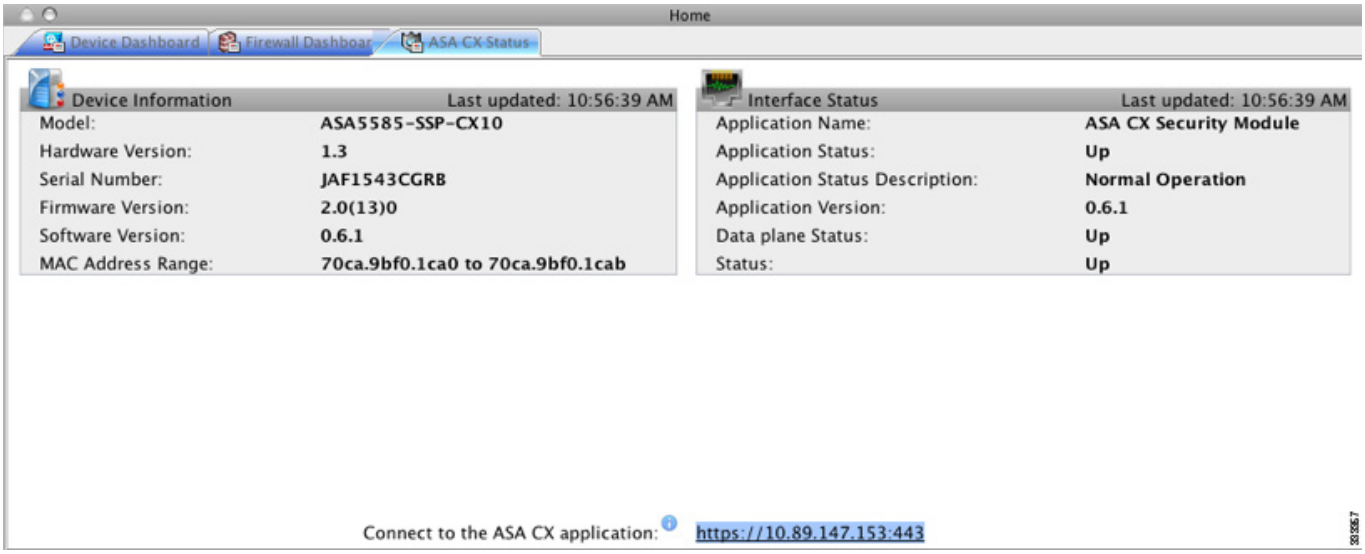


Legend

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.

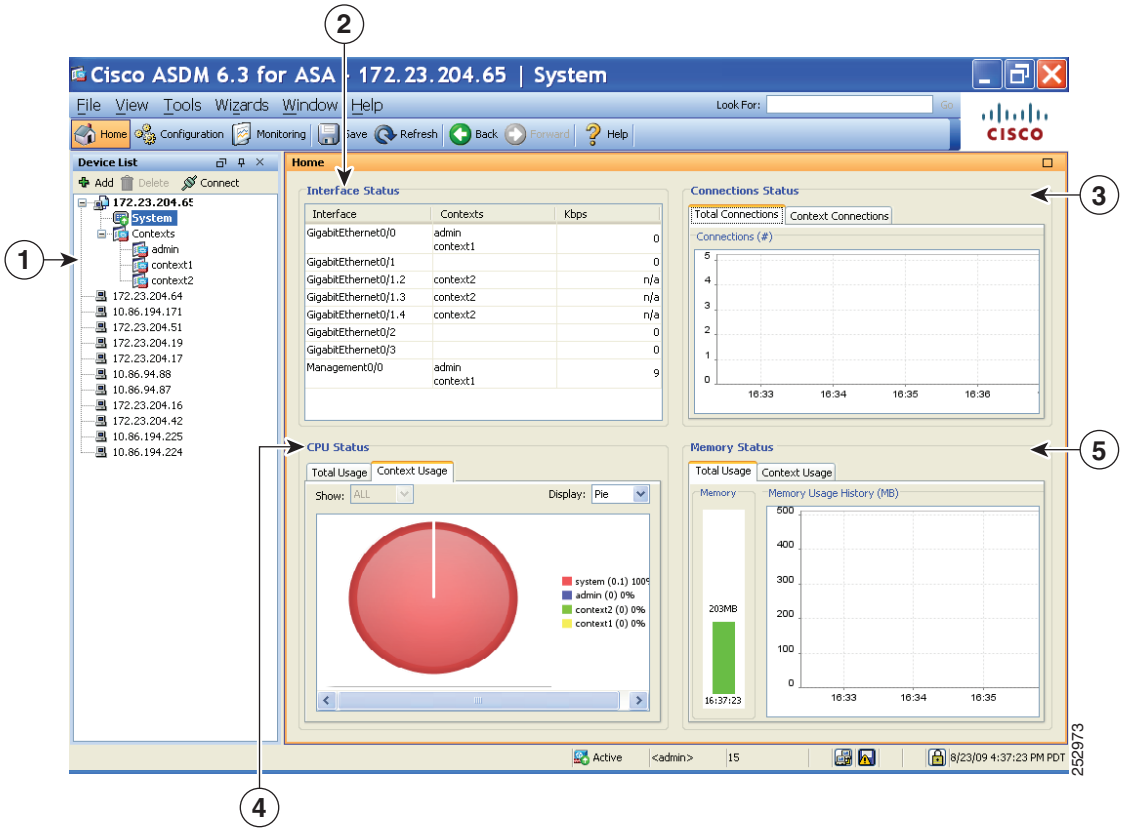


Home Pane (System)

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 9.1; 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 5-39
- 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 5-3
- ASA 5510, page 5-4
- ASA 5520, page 5-5

- ASA 5540, page 5-6
- ASA 5550, page 5-7
- ASA 5580, page 5-8
- ASA 5512-X, page 5-9
- ASA 5515-X, page 5-10
- ASA 5525-X, page 5-11
- ASA 5545-X, page 5-12
- ASA 5555-X, page 5-13
- ASA 5585-X with SSP-10, page 5-14
- ASA 5585-X with SSP-20, page 5-15
- ASA 5585-X with SSP-40 and -60, page 5-16
- ASA Services Module, page 5-17

Items that are in *italics* are separate, optional licenses that can replace the Base or Security Plus license version. 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.

Note Some features are incompatible with each other. See the individual feature chapters for compatibility information.

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.

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, Total UC 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
General Licenses		
Encryption	Base (DES) / *Opt. lic.: Strong (3DES/AES)*	Base (DES) / *Opt. lic.: Strong (3DES/AES)*
Failover	No support	Active/Standby (no stateful failover)
Security Contexts	No support	No support
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, Total UC Proxy Sessions	2 *Optional licenses:* *24* *50* *100*	2 *Optional licenses:* *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*
AnyConnect Premium (sessions)	2 *Optional Perm. or Time-based lic,:* *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*	2 *Optional Perm. or Time-based lic:* *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	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.	364	564
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, Total UC 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.	764								
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, Total UC 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									
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.	964									
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, Total UC 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.	1764									
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, Total UC 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.	4612										
Security Contexts	2	*Optional licenses:*			*5*	*10*	*20*	*50*	*100*	*250*	
Clustering	Disabled	*Optional license: Available for 8 units*									
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

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, Total UC Proxy Sessions	2	*Optional licenses:*	2	*Optional licenses:*
	24	*50*	*100*	*250*	*500*	*24*	*50*	*100*	*250*	*500*
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.	716	916
Security Contexts	No support	2	*Optional licenses:*	*5*
Clustering	No Support	Supported for 2 units
IPS Module	Disabled	*Optional license: Available*	Disabled	*Optional license: Available*
VLANs, Maximum	50	100

ASA 5515-X

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, Total UC 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.	916						
Security Contexts	2	*Optional licenses:*	*5*				
Clustering	Supported for 2 units						
IPS Module	Disabled	*Optional license: Available*					
VLANs, Maximum	100						

ASA 5525-X

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, Total UC 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.	1316								
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*				
Clustering	Supported for 2 units								
IPS Module	Disabled	*Optional license: Available*							
VLANs, Maximum	200								

ASA 5545-X

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, Total UC 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.	1716									
Security Contexts	2	*Optional licenses:*	*5*	*10*	*20*	*50*				
Clustering	Supported for 2 units									
IPS Module	Disabled	*Optional license: Available*								
VLANs, Maximum	300									

ASA 5555-X

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, Total UC 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.	2516									
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*		
Clustering	Supported for 2 units									
IPS Module	Disabled	*Optional license: Available*								
VLANs, Maximum	500									

ASA 5585-X with SSP-10

You can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-10 with an SSP-20 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.

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

Licenses	Description (Base and Security Plus 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, Total UC 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										
10 GE I/O	**Base License**: Disabled; fiber ifcs run at 1 GE					**Security Plus License**: 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.	4612									
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*		
Clustering	Disabled	*Optional license: Available for 8 units*								
VLANs, Maximum	1024									

ASA 5585-X with SSP-20

You can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-20 with an SSP-40 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.

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

Licenses	Description (Base and Security Plus License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	*Optional Time-based license: Available*									
Firewall Conns, Concurrent	2,000,000										
GTP/GPRS	Disabled	*Optional license: Available*									
Intercompany Media Eng.	Disabled	*Optional license: Available*									
UC Phone Proxy Sessions, Total UC 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*									
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	**Base License**: Disabled; fiber ifcs run at 1 GE					**Security Plus License**: 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.	4612										
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*	*250*		
Clustering	Disabled	*Optional license: Available for 8 units*									
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 5585-X with SSP-40 and -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.

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, Total UC 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*									
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.	4612										
Security Contexts	2	*Optional licenses:*		*5*	*10*	*20*	*50*	*100*	*250*		
Clustering	Disabled	*Optional license: Available for 8 units*									
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

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, Total UC 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*									
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										
Security Contexts	2	*Optional licenses:*									
		5	*10*	*20*	*50*	*100*	*250*				
Clustering	No support										
VLANs, Maximum	1000										

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 strong 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. Every **interface** defined in the configuration counts against this limit.
IPS module	The IPS module license lets you run the IPS software module on the ASA. You also need the IPS signature subscription on the IPS side. See the following guidelines: • To buy the IPS signature subscription you need to have the ASA with IPS pre-installed (the part number must include "IPS", for example ASA5515-IPS-K9); you cannot buy the IPS signature subscription for a non-IPS part number ASA. • For failover, you need the IPS signature subscription on both units; this subscription is not shared in failover, because it is not an ASA license. • For failover, the IPS signature subscription requires a unique IPS module license per unit. Like other ASA licenses, the IPS module license is technically shared in the failover cluster license. However, because of the IPS signature subscription requirements, you must buy a separate IPS module license for each unit in failover.

Table 5-16 License Notes (continued)

License	Notes
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.
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, Total UC 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.
VLANs, Maximum	For an interface to count against the VLAN limit, you must assign a VLAN to it.
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 5-34

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 5-39 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 5-39.

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.

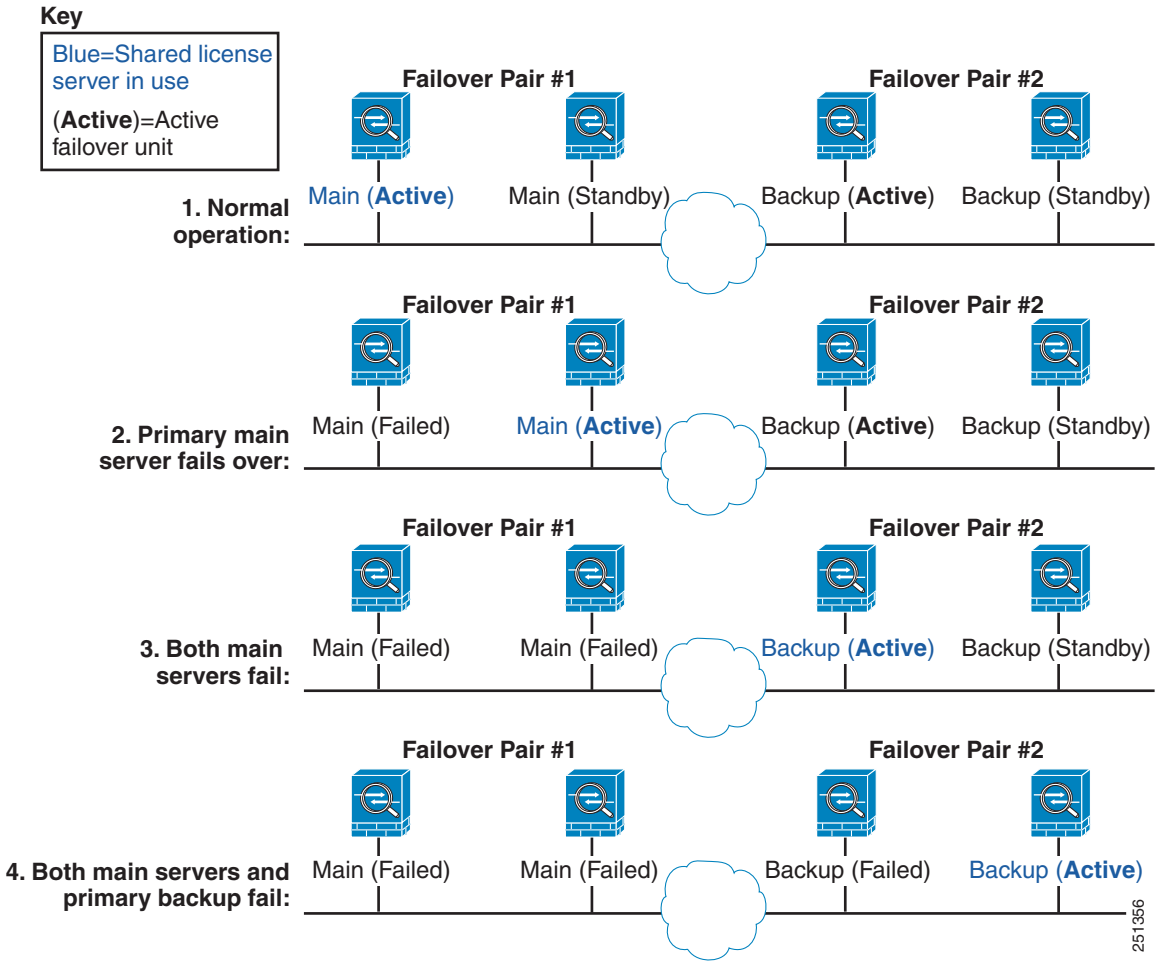
Note The backup server mechanism is separate from, but compatible with, failover.

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 5-33
- 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—Both units require the IPS module license. You also need the IPS signature subscription on the IPS side for both units. See the following guidelines:
  - To buy the IPS signature subscription you need to have the ASA with IPS pre-installed (the part number must include “IPS”, for example ASA5515-IPS-K9); you cannot buy the IPS signature subscription for a non-IPS part number ASA.
  - You need the IPS signature subscription on both units; this subscription is not shared in failover, because it is not an ASA license.
  - The IPS signature subscription requires a unique IPS module license per unit. Like other ASA licenses, the IPS module license is technically shared in the failover cluster license. However, because of the IPS signature subscription requirements, you must buy a separate IPS module license for each unit in.



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 four ASA 5585-X ASAs with SSP-10, three units with 50 contexts each, and one unit with the default 2 contexts. Because the platform limit is 100, the combined license allows a maximum of 100 contexts. Therefore, you can configure up to 100 contexts on the master unit; each slave unit will also have 100 contexts through configuration replication.
  - You have four ASA 5585-X ASAs with SSP-60, three units with 50 contexts each, and one unit with the default 2 contexts. Because the platform limit is 250, the licenses will be combined for a total of 152 contexts. Therefore, you can configure up to 152 contexts on the master unit; each slave unit will also have 152 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 5-39.

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 Filter (which uses SSL).

When you view the license (see the "Monitoring Licenses" section on page 5-39), 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.
- See the "Failover or ASA Cluster Licenses" section on page 5-31.

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 Authorization 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
All models	Downgrading the Encryption license.

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.

Table 5-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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	9.0(1)	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.
ASA 5500-X support for clustering	9.1(4)	The ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X now support 2-unit clusters. Clustering for 2 units is enabled by default in the base license; for the ASA 5512-X, you need the Security Plus license. We did not modify any ASDM screens.



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.

- Information About the Firewall Mode, page 6-1
- Licensing Requirements for the Firewall Mode, page 6-7
- Default Settings, page 6-7
- Guidelines and Limitations, page 6-8
- Setting the Firewall Mode (Single Mode), page 6-9
- Configuring ARP Inspection for the Transparent Firewall, page 6-10
- Customizing the MAC Address Table for the Transparent Firewall, page 6-12
- Firewall Mode Examples, page 6-13
- Feature History for the Firewall Mode, page 6-24

Information About the Firewall Mode

- Information About Routed Firewall Mode, page 6-1
- 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.

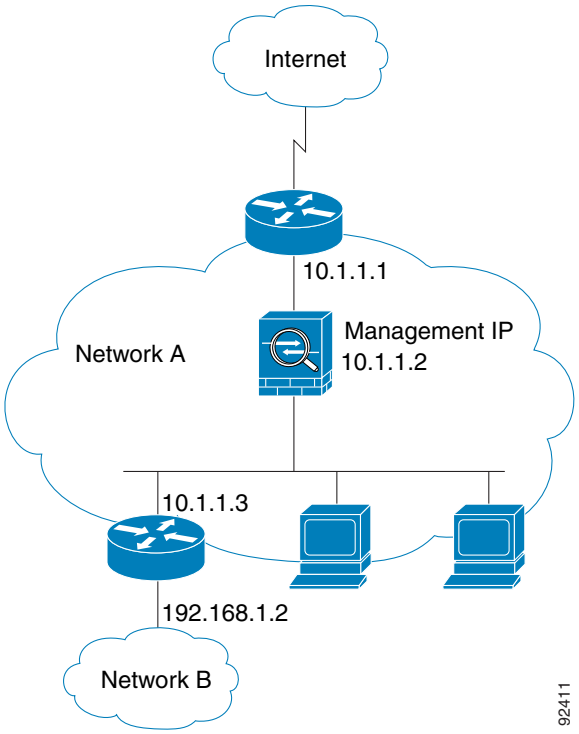
- Using the Transparent Firewall in Your Network, page 6-2
- Bridge Groups, page 6-3
- Management Interface (ASA 5510 and Higher), page 6-4
- Allowing Layer 3 Traffic, page 6-4
- Allowed MAC Addresses, page 6-5
- Passing Traffic Not Allowed in Routed Mode, page 6-5
- BPDU Handling, page 6-5
- MAC Address vs. Route Lookups, page 6-6
- ARP Inspection, page 6-6
- MAC Address Table, page 6-7

Using the Transparent Firewall in Your 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.

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

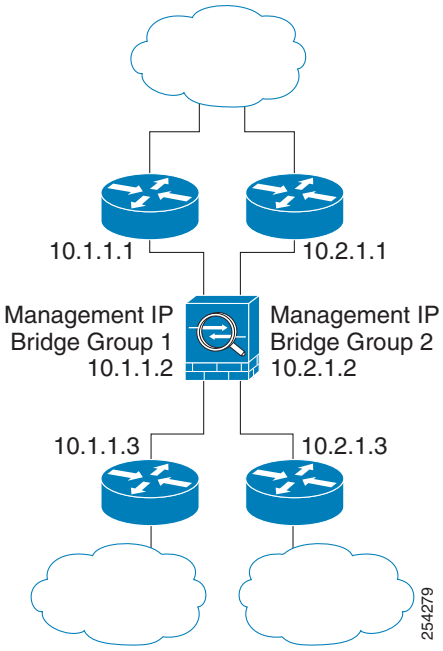


Bridge Groups

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.

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



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 6-4.

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 11-2.

Allowing Layer 3 Traffic

- Unicast IPv4 and IPv6 traffic is allowed through the transparent firewall automatically from a higher security interface to a lower security interface, without an ACL.

Note Broadcast and multicast traffic can be passed using access rules. See the Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules section in the firewall configuration guide for more information.

- ARPs are allowed through the transparent firewall in both directions without an ACL. ARP traffic can be controlled by ARP inspection.
- For Layer 3 traffic travelling from a low to a high security interface, an extended ACL is required on the low security interface. See Chapter 6, “Configuring Access Rules,” in the firewall configuration guide 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 ACL. The transparent firewall, however, can allow almost any traffic through using either an extended ACL (for IP traffic) or an EtherType ACL (for non-IP traffic).

Non-IP traffic (for example AppleTalk, IPX, BPDUs, and MPLS) can be configured to go through using an EtherType ACL.

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. An exception is made for BPDUs and IS-IS, 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 ACL, 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 ACL. 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 ACL 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 9-14 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 to find the next hop gateway; you need to add a static route on the ASA for the real host address.
- Voice over IP (VoIP) and DNS 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. If you enable NAT for the inspected traffic, a static route is required to determine the egress interface for the real host address that is embedded in the packet. Affected applications include:
  - CTIQBE
  - DNS
  - GTP
  - H.323
  - MGCP
  - RTSP
  - SIP
  - Skinny (SCCP)

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.

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 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.

Transparent Mode Defaults

- By default, all ARP packets are allowed through the ASA.
- If you enable ARP inspection, the default setting is to flood non-matching packets.
- 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

Set the firewall mode per context.

Transparent Firewall 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.
- 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 14-5 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 6-9 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 ACLs: 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 ACL.
Multicast IP routing	You can allow multicast traffic through the ASA by allowing it in an extended ACL.
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 ACL, but it does not terminate non-management connections. Clientless 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 and for the currently connected context in multiple mode (typically the admin context), you cannot change the mode in ASDM. For other multiple mode contexts, you can set the mode in ASDM for each context; see the "Configuring a Security Context" section on page 8-20.

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 6-8 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

Note To set the firewall mode to transparent and also configure ASDM management access after the configuration is cleared, see the "Customizing ASDM Access (ASA 5505)" section on page 3-8 or "Customizing ASDM Access (ASA 5510 and Higher)" section on page 3-10.

Command	Purpose

`firewall transparent`
 	
Example:
`ciscoasa(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.

Configuring ARP Inspection for the Transparent Firewall

This section describes how to configure ARP inspection and includes the following topics:

- Task Flow for Configuring ARP Inspection, page 6-10
- Adding a Static ARP Entry, page 6-10
- Enabling ARP Inspection, page 6-11

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 6-10. 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 6-11.

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 6-11.

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**.

Customizing the MAC Address Table for the Transparent Firewall

This section describes how you can customize the MAC address table and includes the following sections:

- Adding a Static MAC Address, page 6-12
- Disabling MAC Address Learning, page 6-13

Adding a Static MAC Address

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can add static MAC addresses to the MAC address table if desired. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the 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 6-10), 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**.

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 6-13
- How Data Moves Through the Transparent Firewall, page 6-19

How Data Moves Through the ASA in Routed Firewall Mode

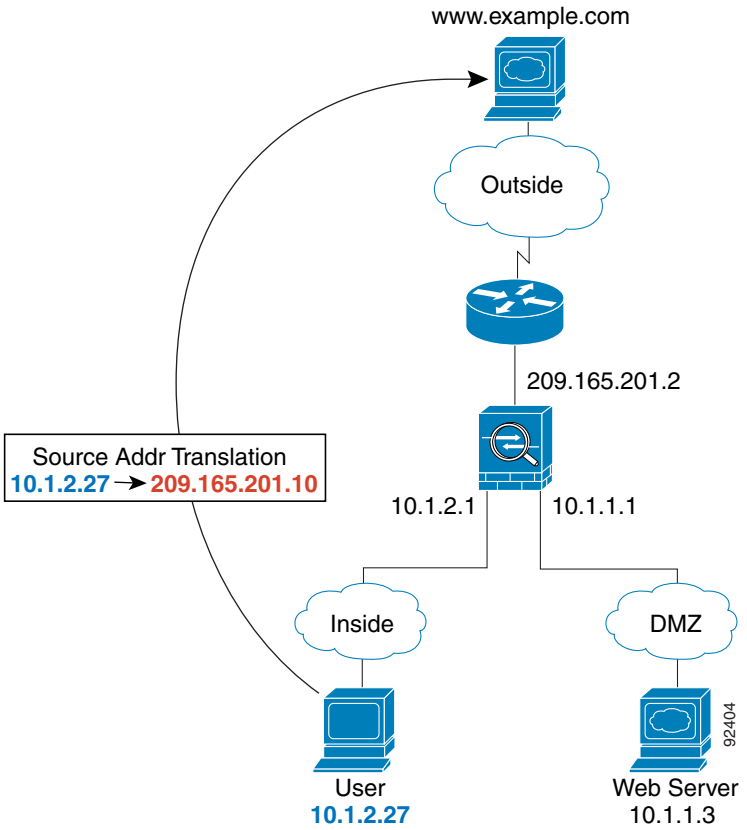
This section describes how data moves through the ASA in routed firewall mode and includes the following topics:

- An Inside User Visits a Web Server, page 6-14
- An Outside User Visits a Web Server on the DMZ, page 6-15
- An Inside User Visits a Web Server on the DMZ, page 6-16
- An Outside User Attempts to Access an Inside Host, page 6-16
- A DMZ User Attempts to Access an Inside Host, page 6-18

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).

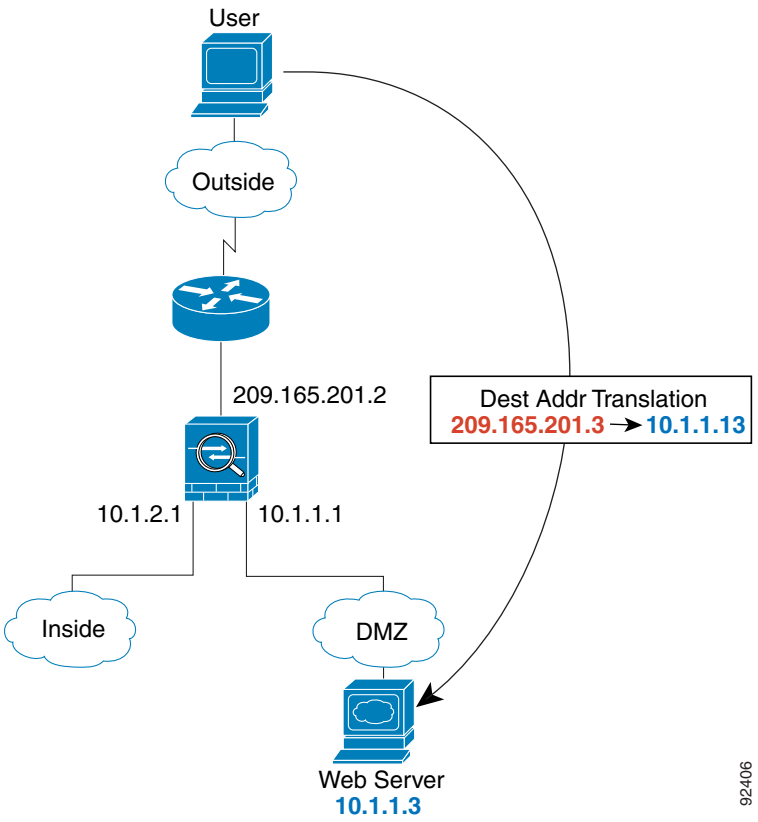
   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA translates the local source address (10.1.2.27) to the global address 209.165.201.10, which is on the outside interface subnet.

   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):

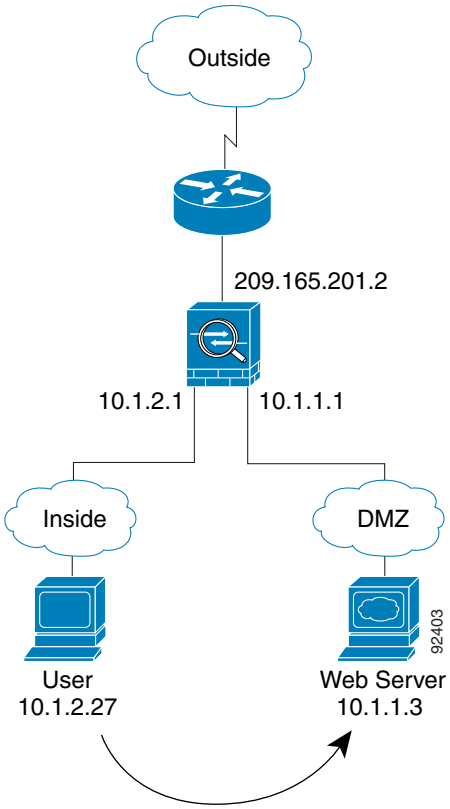
1. A user on the outside network requests a web page from the DMZ web server using the global destination address of 209.165.201.3, which is on the outside interface subnet.
2. The ASA receives the packet and untranslates the destination address to the local address 10.1.1.3.
3. Because it is a new session, 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):

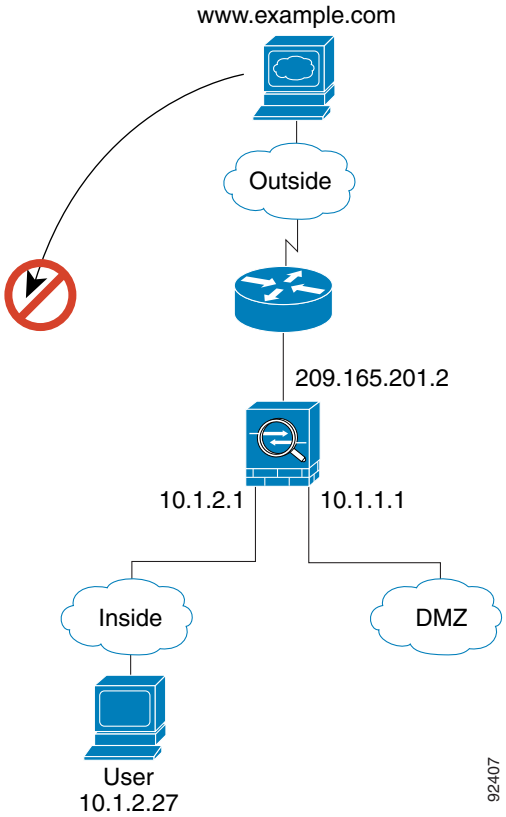
1. A user on the inside network requests a web page from the DMZ web server using the destination address of 10.1.1.3.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   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).

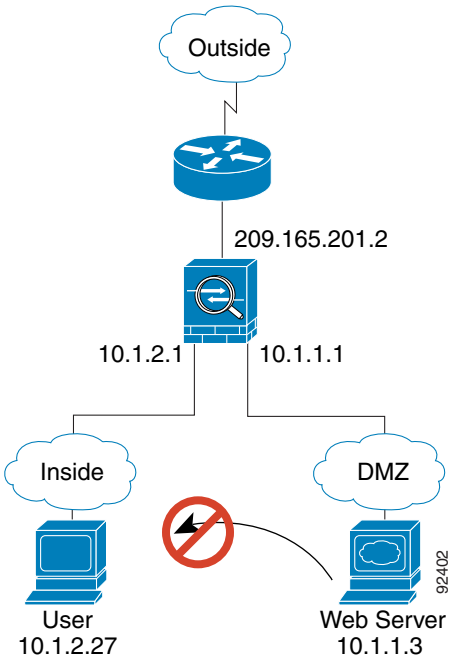
3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

   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):

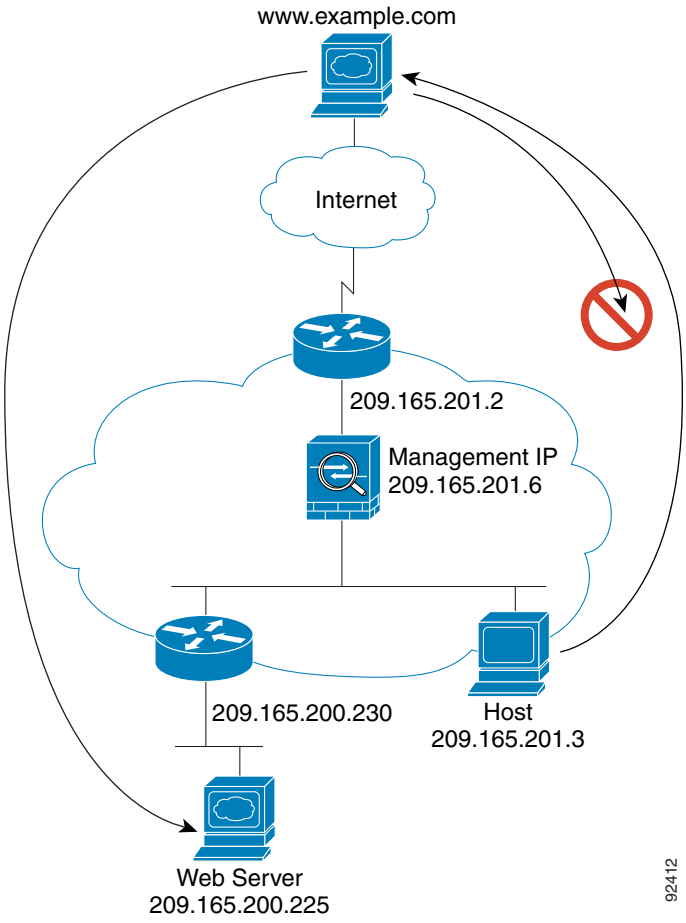
1. A user on the DMZ network attempts to reach an inside host. Because the DMZ does not have to route the traffic on the Internet, the private addressing scheme does not prevent routing.
2. The ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according 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



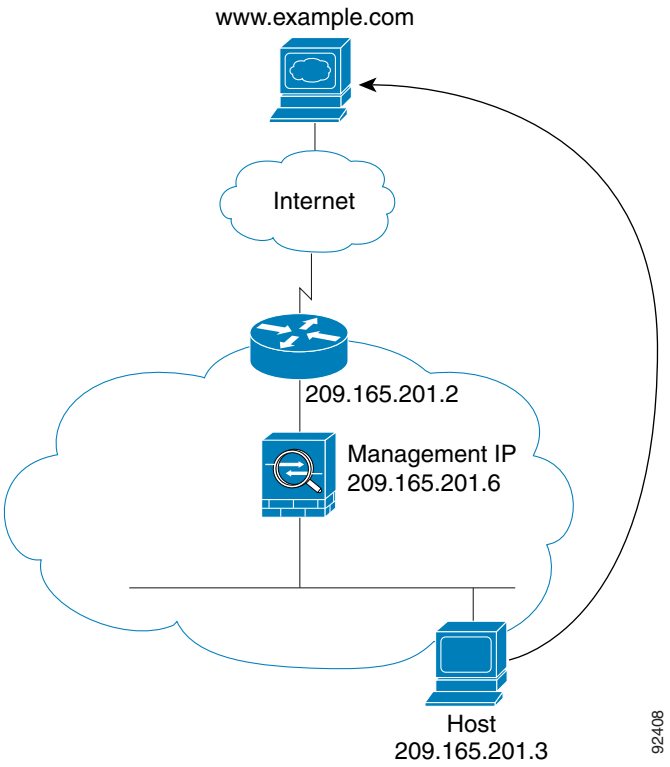
This section describes how data moves through the ASA and includes the following topics:

- An Inside User Visits a Web Server, page 6-20
- An Inside User Visits a Web Server Using NAT, page 6-21
- An Outside User Visits a Web Server on the Inside Network, page 6-22
- An Outside User Attempts to Access an Inside Host, page 6-23

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).

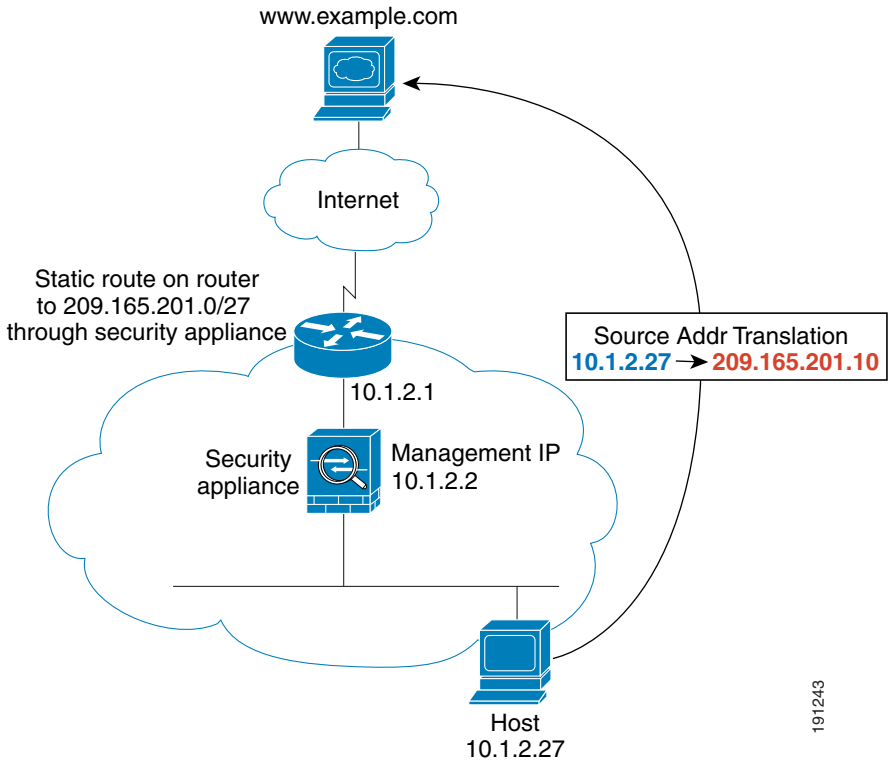
   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 209.165.201.2.

   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.

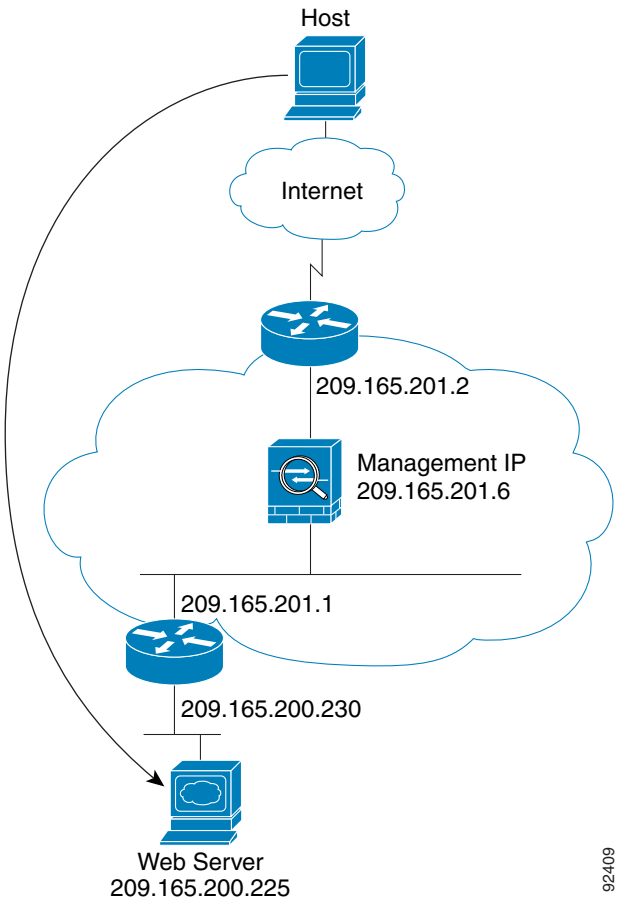
   Because the mapped address is not on the same network as the outside interface, then be sure the upstream router has a static route to the mapped network that points to the ASA.
4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. If the destination 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).

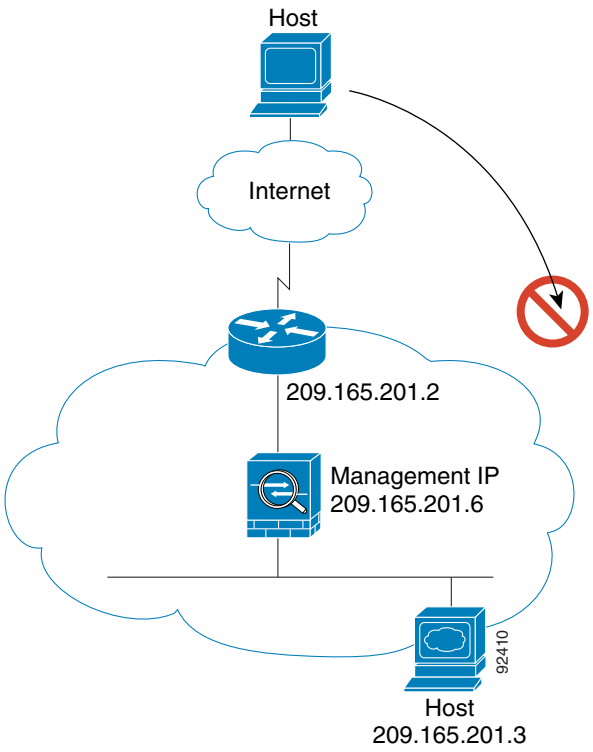
   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.

   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.

Feature History for the Firewall Mode

Table 6-2 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 6-2 Feature History for Firewall Mode

Feature Name	Platform 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.
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**.
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**.
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

Table 6-2 Feature History for Firewall Mode (continued)

Feature Name	Platform Releases	Feature Information
ARP cache additions for non-connected subnets	8.4(5)/9.1(2)	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.
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.



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
- ASA CX Basic Configuration (ASA 5585-X), 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 16-1.

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 12-6 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 12-6 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 13-7 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 13-11 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 13-7 for more information..
- To configure the IPv6 address, see the "Configuring IPv6 Addressing" section on page 13-15.

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 13-11 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 14-8 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 13-7 for more information.
- See the "Allowing Same Security Level Communication" section on page 13-20 for information about the Enable traffic between... check boxes.

Static Routes

Configure static routes. See Chapter 25, "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 in the VPN configuration guide.

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 the DHCP Server" section on page 17-4 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 section or the Configuring Dynamic PAT (Hide) section in the firewall configuration guide for more information.

Administrative Access

- Configures ASDM, Telnet, or SSH access. See the "Configuring Management Access" section on page 45-3 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 45-3 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 in the firewall configuration guide.

ASA CX Basic Configuration (ASA 5585-X)

You can use the Startup Wizard in ASDM to configure the ASA CX management address and Auth Proxy Port. These settings are saved to the ASA CX configuration, not the ASA configuration. **Note**: You will also need to set additional network settings at the ASA CX CLI. See the (ASA 5585-X) Changing the ASA CX Management IP Address section in the firewall configuration guide for information about this screen.

Time Zone and Clock Configuration

Configure the clock parameters. See the "Setting the Date and Time" section on page 16-3 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 46-31 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 2

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 8-1
- Licensing Requirements for Multiple Context Mode, page 8-13
- Guidelines and Limitations, page 8-14
- Default Settings, page 8-15
- Configuring Multiple Contexts, page 8-15
- Changing Between Contexts and the System Execution Space, page 8-25
- Managing Security Contexts, page 8-26
- Monitoring Security Contexts, page 8-31
- Feature History for Multiple Context Mode, page 8-34

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 8-14.

This section provides an overview of security contexts and includes the following topics:

- Common Uses for Security Contexts, page 8-2
- Context Configuration Files, page 8-2
- How the ASA Classifies Packets, page 8-3
- Cascading Security Contexts, page 8-6
- Management Access to Security Contexts, page 8-7
- Information About Resource Management, page 8-8
- Information About MAC Addresses, page 8-11

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 8-2
- System Configuration, page 8-2
- Admin Context Configuration, page 8-2

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 8-3
- Classification Examples, page 8-4

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 8-3
- Unique MAC Addresses, page 8-3
- NAT Configuration, page 8-3

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.

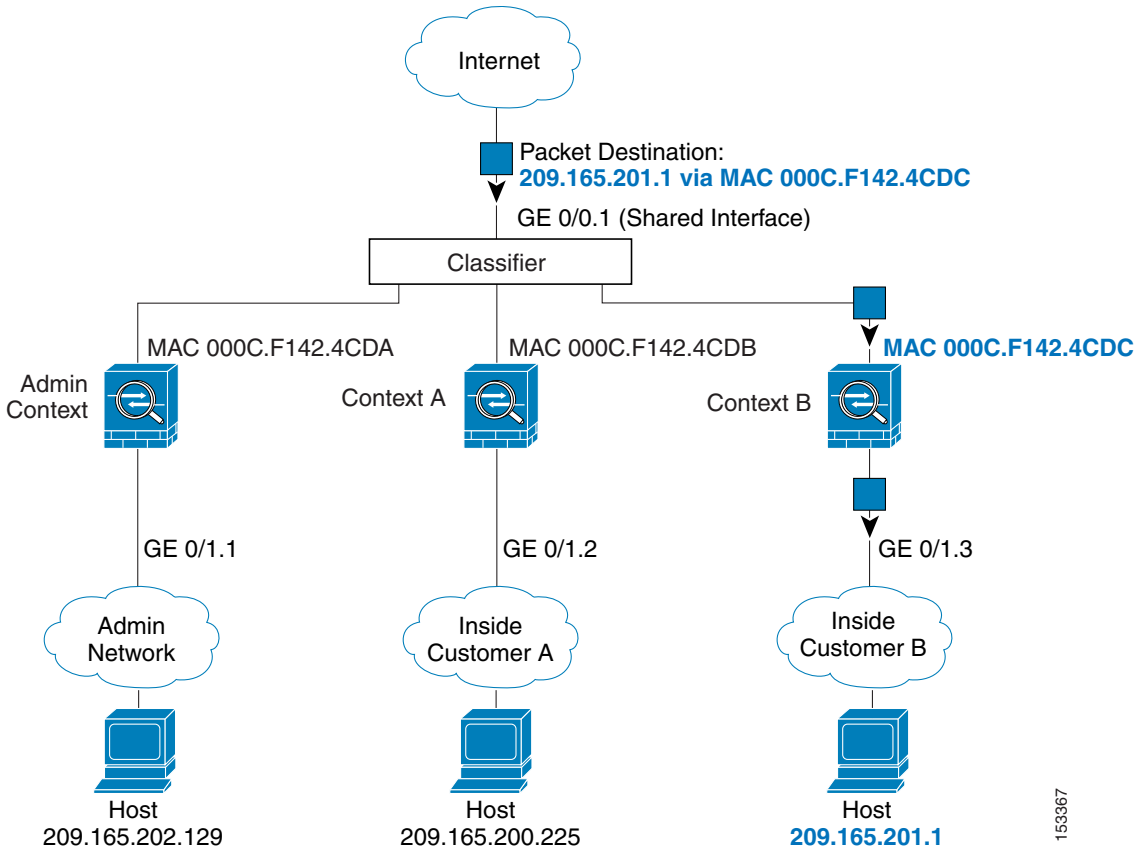
NAT Configuration

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

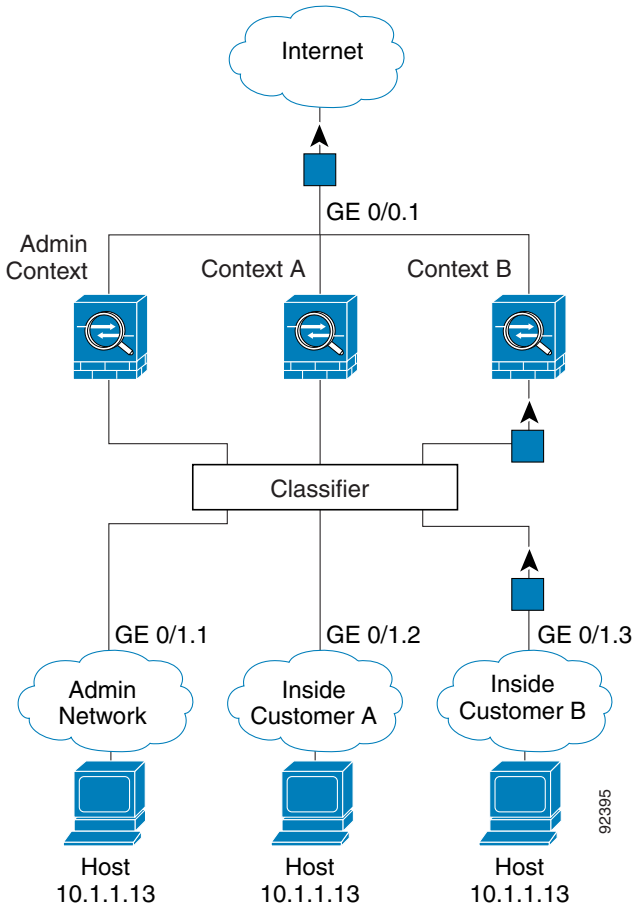
Figure 8-1 shows multiple contexts sharing an outside interface. The classifier assigns the packet to Context B because Context B includes the MAC address to which the router sends the packet.

Figure 8-1 Packet Classification with a Shared Interface Using MAC Addresses



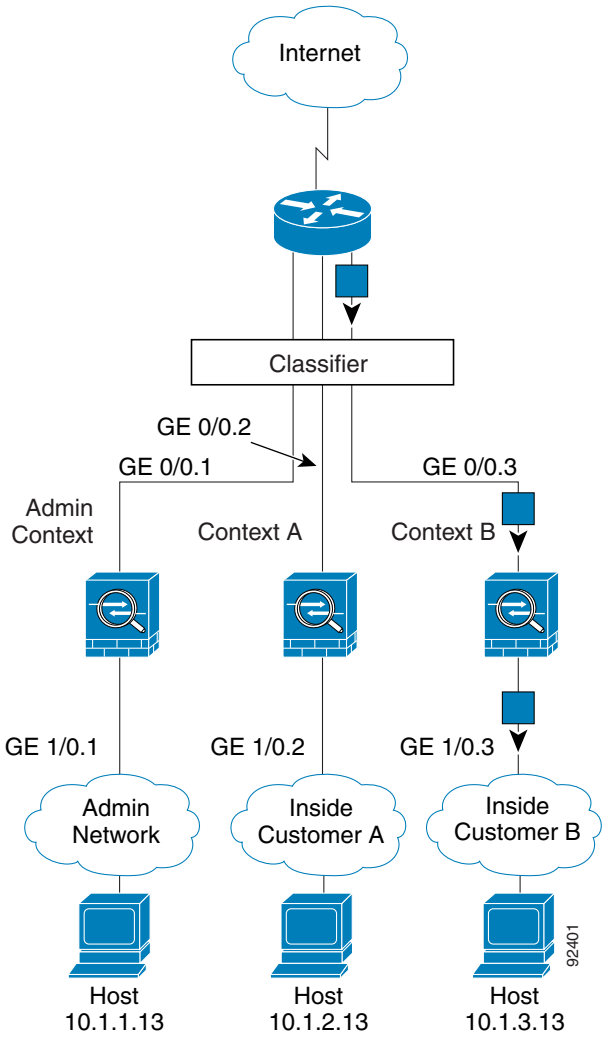
Note that all new incoming traffic must be classified, even from inside networks. Figure 8-2 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.

Figure 8-2 Incoming Traffic from Inside Networks



For transparent firewalls, you must use unique interfaces. Figure 8-3 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 8-3 Transparent Firewall Contexts



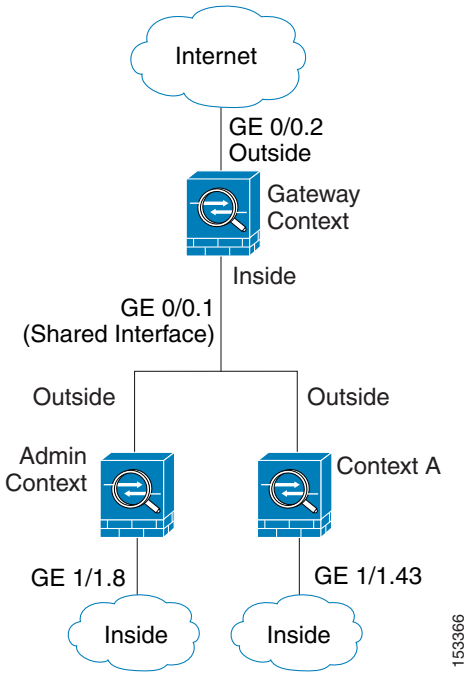
Cascading Security Contexts

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 8-4 shows a gateway context with two contexts behind the gateway.

Figure 8-4 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 8-7
- Context Administrator Access, page 8-8

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 45, "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 45, "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 8-8
- Resource Limits, page 8-8
- Default Class, page 8-9
- Using Oversubscribed Resources, page 8-10
- Using Unlimited Resources, page 8-11

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.

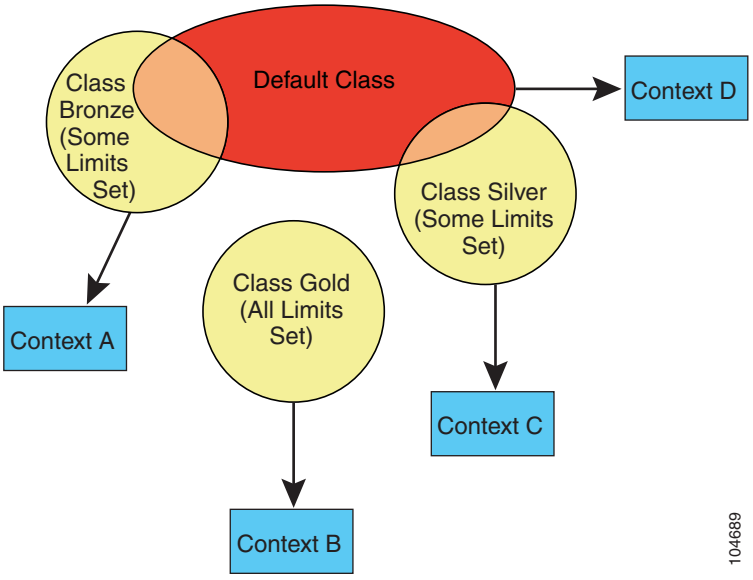
If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 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 8-5 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.

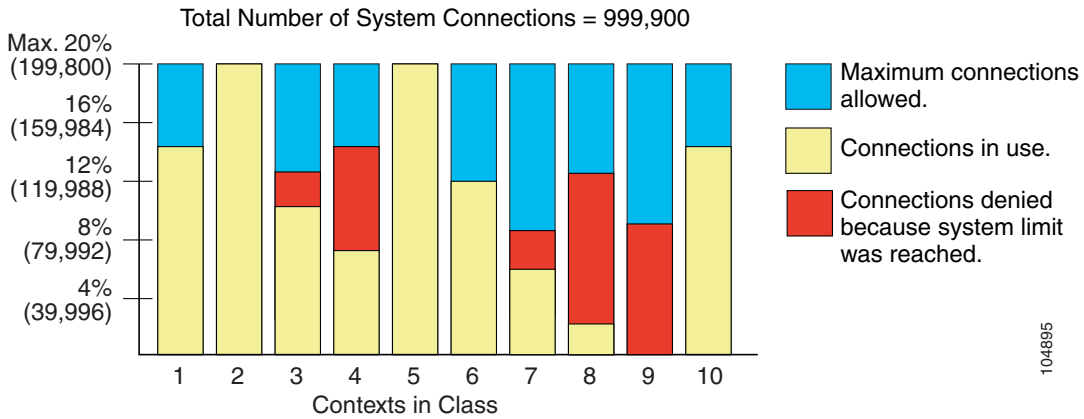
Figure 8-5 Resource Classes



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 8-6.)

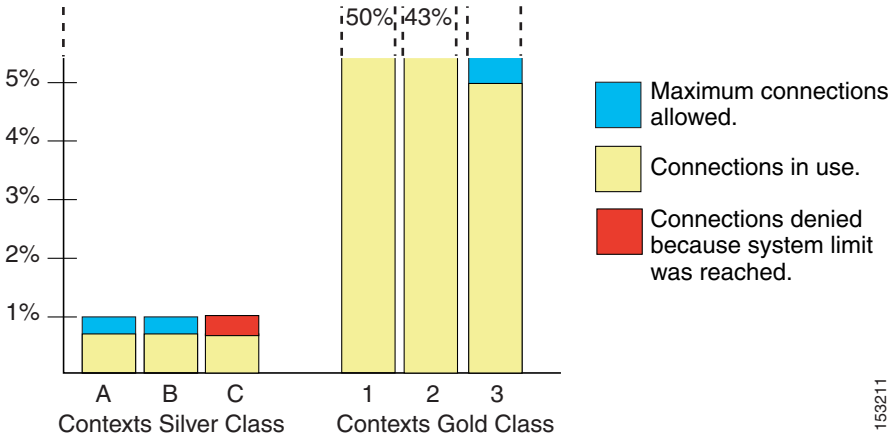
Figure 8-6 Resource Oversubscription



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 8-7.) Setting unlimited access is similar to oversubscribing the ASA, except that you have less control over how much you oversubscribe the system.

Figure 8-7 Unlimited Resources



Information About MAC Addresses

To allow contexts to share interfaces, the ASA assigns virtual MAC addresses to each shared context interface by default. To customize or disable auto-generation, see the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 8-24.

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 other classification methods are attempted that might not provide full coverage. See the "How the ASA Classifies Packets" section on page 8-3 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, MTU, and TCP MSS" section on page 13-13 to manually set the MAC address.

This section includes the following topics:

- Default MAC Address, page 8-12
- Interaction with Manual MAC Addresses, page 8-12
- Failover MAC Addresses, page 8-12
- MAC Address Format, page 8-12

Default MAC Address

(8.5(1.7) and Later) Automatic MAC address generation is enabled by default. The ASA autogenerates the prefix based on the last two bytes of the interface (ASA 5500) or backplane (ASASM) MAC address. You can customize the prefix if desired.

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 also the "MAC Address Format" section on page 8-12.

Note (8.5(1.6) and earlier) To maintain hitless upgrade for failover pairs, the ASA does not convert an existing legacy 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. For more information about the legacy method, see the **mac-address auto** command in the command reference.

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 8-12 section for more information.

MAC Address Format

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*

Note The MAC address format without a prefix is a legacy version not supported on newer ASA versions. See the **mac-address auto** command in the command reference for more information about the legacy format.

Licensing Requirements for Multiple Context Mode

Model	License Requirement
ASA 5505	No support.
ASA 5510	• Base License: No support. • 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.*
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	• Base License: No support. • Security Plus License: 2 contexts. *Optional license: 5 contexts.*
ASA 5515-X	Base 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.*

Model	License Requirement
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.

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.
- If you store context configurations in the root directory of flash memory, on some models you might run out of room in that directory, even though there is available memory. In this case, create a subdirectory for your configuration files. Background: some models, such as the ASA 5585-X, use the FAT 16 file system for internal flash memory, and if you do not use 8.3-compliant short names, or use uppercase characters, then fewer than 512 files and folders can be stored because the file system uses up slots to store long file names (see http://support.microsoft.com/kb/120138/en-us).

Default Settings

- By default, the ASA is in single context mode.
- See the "Default Class" section on page 8-9.
- See the "Default MAC Address" section on page 8-12.

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 8-15
- Enabling or Disabling Multiple Context Mode, page 8-16
- Configuring a Class for Resource Management, page 8-17
- Configuring a Security Context, page 8-20
- Automatically Assigning MAC Addresses to Context Interfaces, page 8-24

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 8-16.

Step 2 (Optional) Configure classes for resource management. See the "Configuring a Class for Resource Management" section on page 8-17. **Note**: For VPN support, you must configure VPN resources in a resource class; the default class does not allow VPN.

Step 3 Configure interfaces in the system execution space.

- ASA 5500—Chapter 11, "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 8-20.

Step 5 (Optional) Customize MAC address assignments. See the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 8-24.

Step 6 Complete interface configuration in the context. See Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 9, "Configuring 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 8-16
- Restoring Single Context Mode, page 8-16

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 46-12.

Detailed Steps

Command	Purpose

`mode multiple`
 	
Example:
`ciscoasa(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:** `ciscoasa(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:** `ciscoasa(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 8-1 lists the resource types and the limits.

Table 8-1 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[2]	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 8-1 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[2]	Concurrent	N/A	N/A	Network address translations.

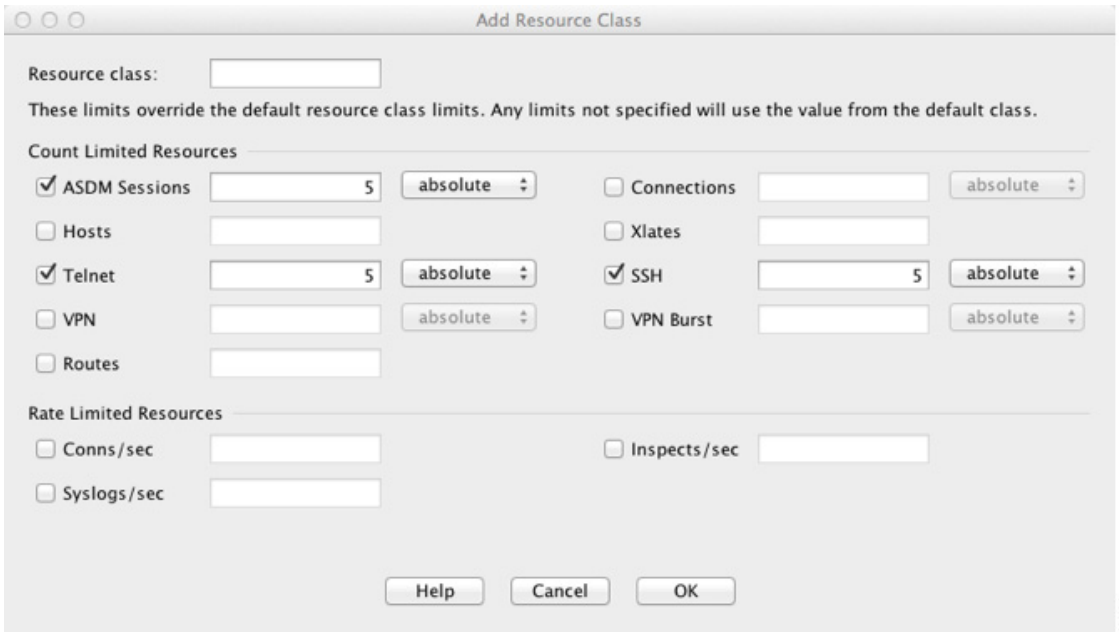
1. If this column value is N/A, then you cannot set a percentage of the resource because there is no hard system limit for the resource.
2. Syslog messages are generated for whichever limit is lower xlates or conns. For example, if you set the xlates limit to 7 and the conns to 9, then the ASA only generates syslog message 321001 (“Resource 'xlates' limit of 7 reached for context 'ctx1'”) and not 321002 (“Resource 'conn rate' limit of 5 reached for context 'ctx1'”).

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 8-1 on page 8-18 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 8-1 on page 8-18 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

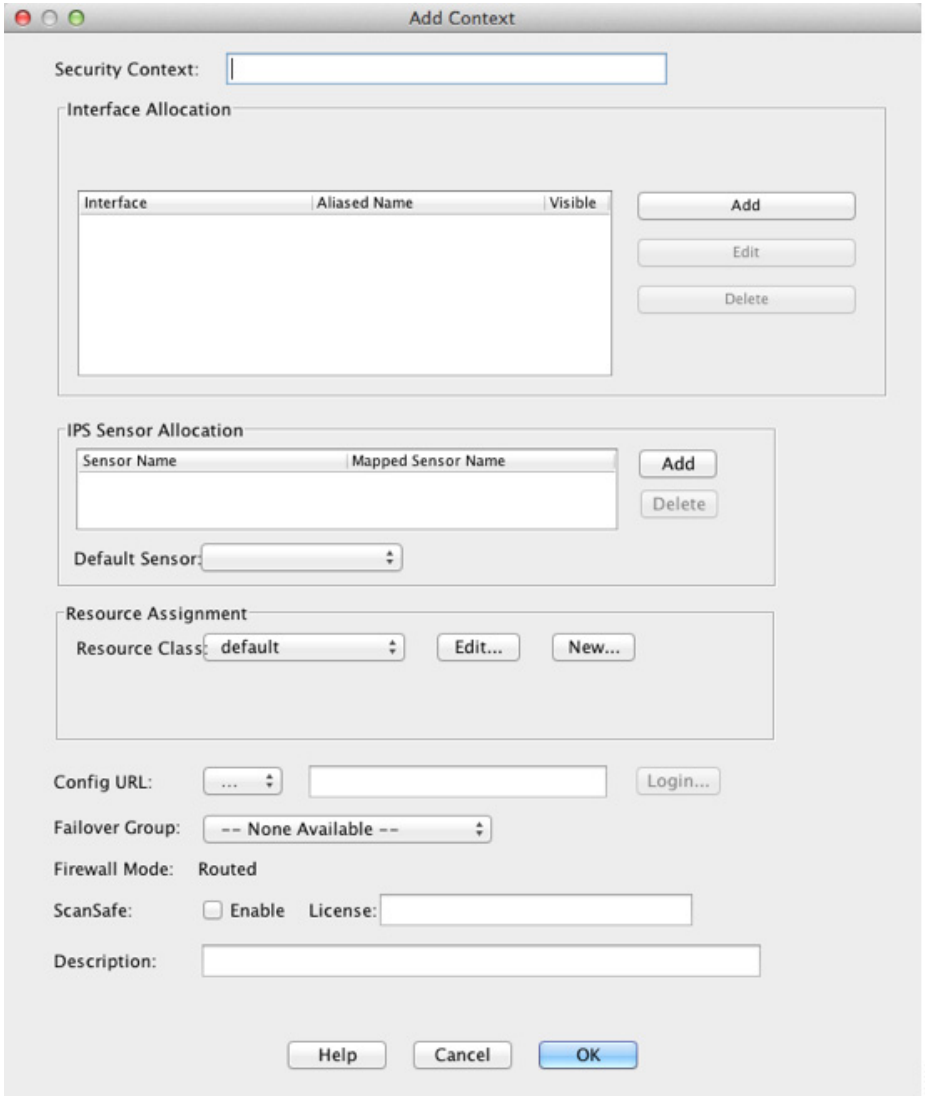
- Perform this procedure in the system execution space.
- For the ASASM, assign VLANs to the ASASM on the switch according to Chapter 2, "Configuring the Switch for Use with the ASA Services Module."
- For the ASA 5500, configure physical interface parameters, VLAN subinterfaces, EtherChannels, and redundant interfaces according to Chapter 11, "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**.

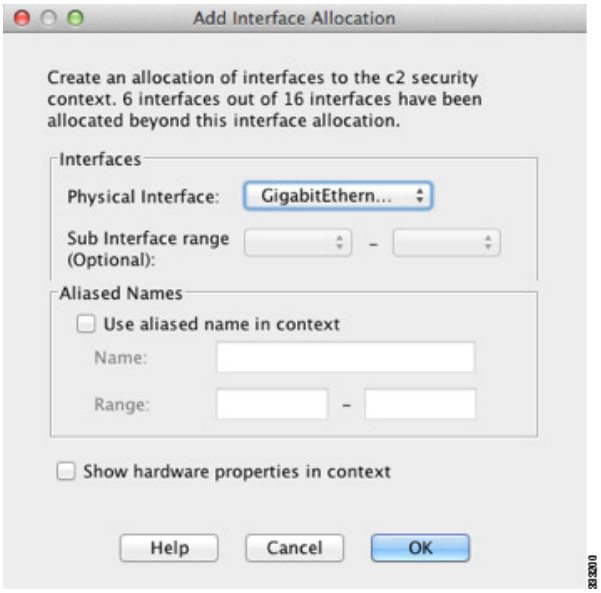
The Add Context dialog box appears.



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 31, “Configuring the ASA IPS Module,” in the firewall configuration guide.

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 8-17 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

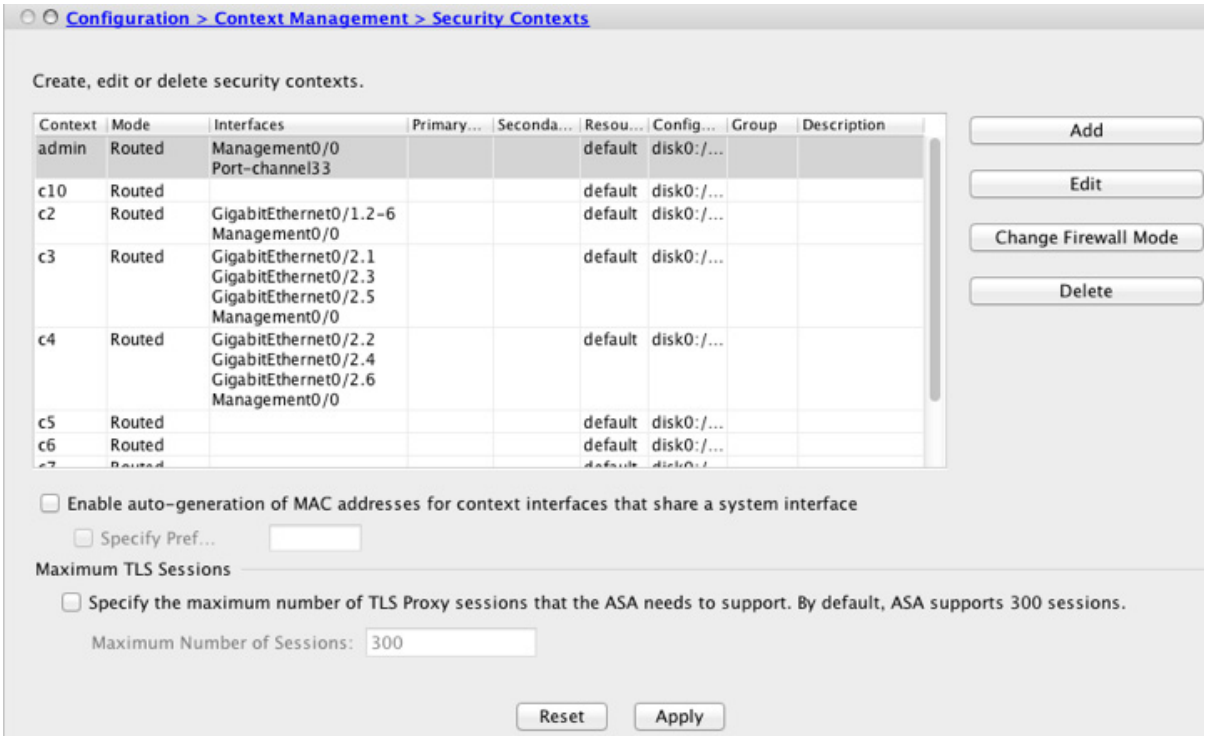
a. (Optional) For external file systems, set the username and password by clicking **Login**.

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.

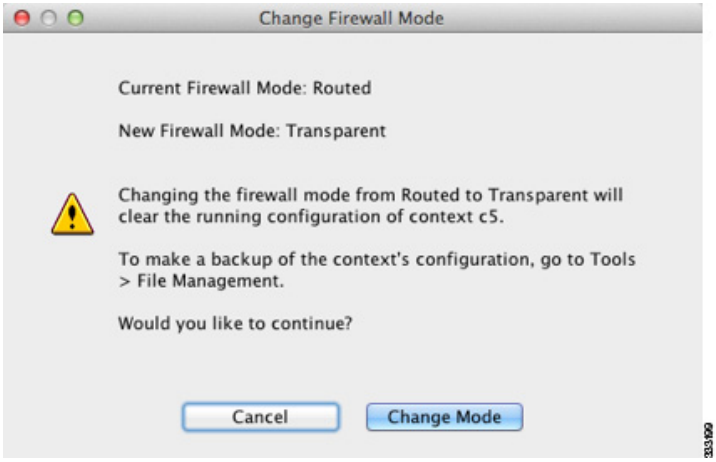
Step 10 (Optional) Add a description in the Description field.

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.

Note You cannot change the mode of the currently connected context in ASDM (typically the admin context); see the "Setting the Firewall Mode (Single Mode)" section on page 6-9 to set the mode at the command line.

Step 13 To customize auto-generation of MAC addresses, see the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 8-24.

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 17, "Configuring the TLS Proxy for Encrypted Voice Inspection," in the firewall configuration guide.

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 8-11 for more information, especially if you are upgrading from an earlier ASA version. See also the "Viewing Assigned MAC Addresses" section on page 8-32.

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, MTU, and TCP MSS" section on page 13-13 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.

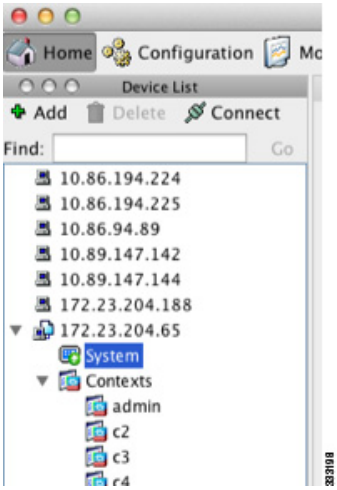
This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address. See the "MAC Address Format" section on page 8-12 section for more information about how the prefix is used.

Changing Between Contexts and the System Execution Space

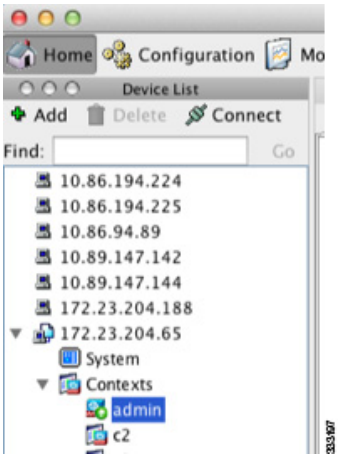
If you log in to the system execution space (or the admin context), you can change between contexts and perform configuration and monitoring tasks within each context. The running configuration that you edit in a configuration mode depends on your location. When you are in the system execution space, 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 8-26
- Changing the Admin Context, page 8-27
- Changing the Security Context URL, page 8-28
- Reloading a Security Context, page 8-29

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

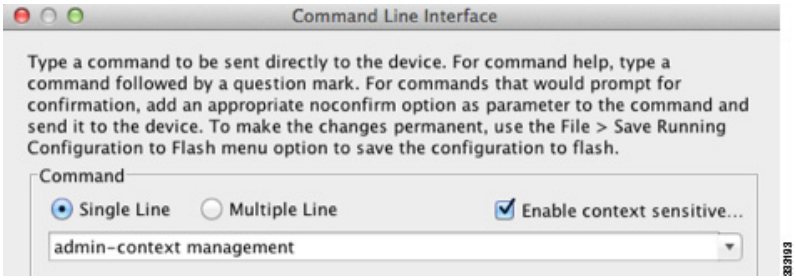
Perform this procedure in the system execution space.

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.

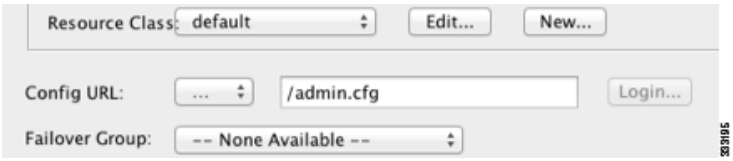
Detailed Steps

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 8-30
- Reloading by Removing and Re-adding the Context, page 8-31

Reloading by Clearing the Configuration

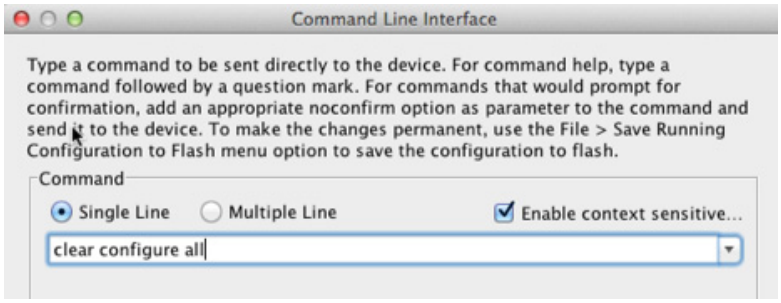
To reload the context by clearing the context configuration and reloading the configuration from the URL, perform the following steps.

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:

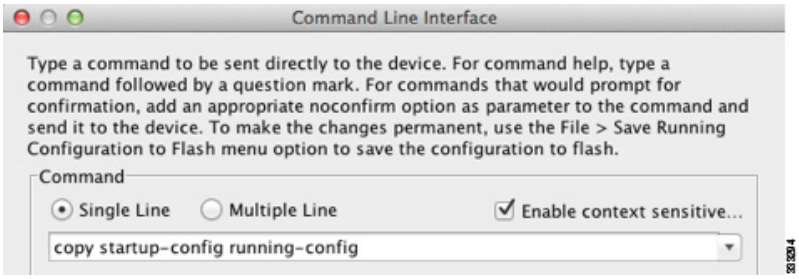
```
clear configure all
```

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 8-26. Be sure to uncheck the **Also delete config URL file from the disk** check box.
2. "Configuring a Security Context" section on page 8-20

Monitoring Security Contexts

This section describes how to view and monitor context information and includes the following topics:

- Monitoring Context Resource Usage, page 8-31
- Viewing Assigned MAC Addresses, page 8-32

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 8-33

- Viewing MAC Addresses Within a Context, page 8-33

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 8-2 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 8-2 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.
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 8-2 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 Failover

This chapter describes how to configure Active/Standby or Active/Active failover, and includes the following sections:

- Introduction to Failover, page 9-1
- Licensing Requirements Failover, page 9-23
- Prerequisites for Failover, page 9-23
- Guidelines and Limitations, page 9-24
- Default Settings, page 9-24
- Configuring Active/Standby Failover, page 9-25
- Configuring Active/Active Failover, page 9-32
- Configuring Optional Failover Parameters, page 9-41
- Managing Failover, page 9-47
- Monitoring Failover, page 9-52
- Feature History for Failover, page 9-54

Introduction to Failover

- Failover Overview, page 9-2
- Failover System Requirements, page 9-2
- Failover and Stateful Failover Links, page 9-3
- MAC Addresses and IP Addresses, page 9-7
- Intra- and Inter-Chassis Module Placement for the ASA Services Module, page 9-8
- Stateless and Stateful Failover, page 9-12
- Transparent Firewall Mode Requirements, page 9-14
- Failover Health Monitoring, page 9-16
- Failover Times, page 9-17
- Configuration Synchronization, page 9-17
- Information About Active/Standby Failover, page 9-19
- Information About Active/Active Failover, page 9-20

Failover Overview

Configuring failover requires two identical ASAs connected to each other through a dedicated failover link and, optionally, a state link. The health of the active units and interfaces is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

The ASA supports two failover modes, Active/Active failover and Active/Standby failover. Each failover mode has its own method for determining and performing failover.

- 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.
- In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into 2 *failover groups*. A failover group is simply a logical group of one or more security contexts. One group is assigned to be active on the primary ASA, and the other group is assigned to be active on the secondary ASA. When a failover occurs, it occurs at the failover group level.

Both failover modes support stateful or stateless failover.

Failover System Requirements

This section describes the hardware, software, and license requirements for ASAs in a failover configuration.

- Hardware Requirements, page 9-2
- Software Requirements, page 9-2
- License Requirements, page 9-3

Hardware Requirements

The two units in a failover configuration must:

- Be the same model.
- Have the same number and types of interfaces.
- Have the same modules installed (if any)
- Have 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 firewall mode (routed or transparent).
- Be in the same context mode (single or multiple).

- Have the same major (first number) and minor (second number) software version. However, you can temporarily 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.

  See the "Upgrading a Failover Pair or ASA Cluster" section on page 46-6 for more information about upgrading the software on a failover pair.
- Have the same AnyConnect images. If the failover pair has mismatched images when a hitless upgrade is performed, then the clientless SSL VPN 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."

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

The failover link and the optional Stateful Failover link are dedicated connections between the two units.

- Failover Link, page 9-3
- Stateful Failover Link, page 9-4
- Avoiding Interrupted Failover and Data Links, page 9-5

Caution All information sent over the failover and state links is sent in clear text unless you secure the communication with an IPsec tunnel or 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 an IPsec tunnel or a failover key if you are using the ASA to terminate VPN tunnels.

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit.

- Failover Link Data, page 9-3
- Interface for the Failover Link, page 9-4
- Connecting the Failover Link, page 9-4

Failover Link Data

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

Interface for the Failover Link

You can use any unused interface (physical, redundant, or EtherChannel) 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 can only be used for the failover link (and optionally also for the state link).

Connecting the 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 an Ethernet cable to connect the units directly, without the need for an external switch.

If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface 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.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link (also known as the state link) to pass connection state information.

You have three interface options for the state link:

- Dedicated Interface (Recommended), page 9-4
- Shared with the Failover Link, page 9-5
- Shared with a Regular Data Interface (Not Recommended), page 9-5

Note Do not use a management interface for the state link.

Dedicated Interface (Recommended)

You can use a dedicated interface (physical, redundant, or EtherChannel) for the state link. 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 an Ethernet cable to connect the appliances directly, without the need for an external switch.

  If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface 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.

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.

Shared with the Failover Link

Sharing a failover link might be necessary if you do not have enough interfaces. If you use the failover link as the state link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the state link.

Shared with a Regular Data Interface (Not Recommended)

Sharing a data interface with the state link 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.

Using a data interface as the state link is supported in single context, routed mode only.

Avoiding Interrupted Failover and Data Links

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the ASA can use the data interfaces to determine if a failover is required. Subsequently, the failover operation is suspended until the health of the failover link is restored.

See the following connection scenarios to design a resilient failover network.

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 9-1 and Figure 9-2 are NOT recommended.

Figure 9-1 Connecting with a Single Switch—Not Recommended

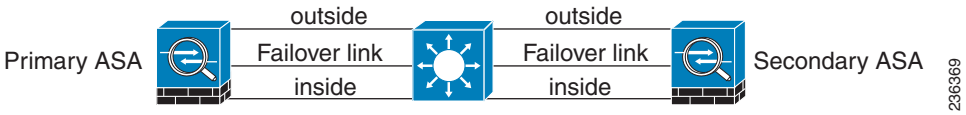
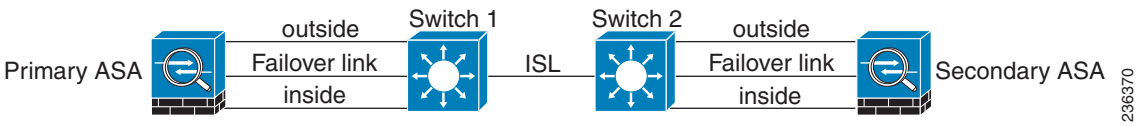


Figure 9-2 Connecting with a Double Switch—Not Recommended



Scenario 2—Recommended

We recommend that failover links NOT use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in Figure 9-3 and Figure 9-4.

Figure 9-3 Connecting with a Different Switch

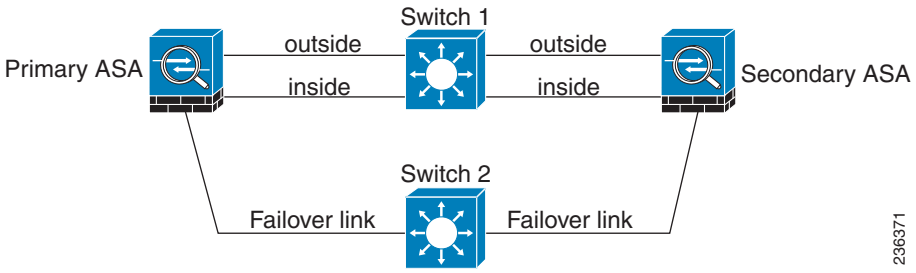
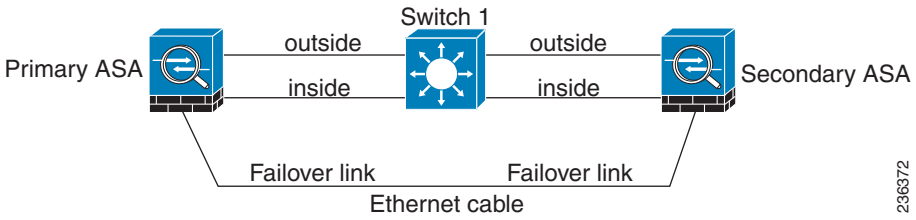


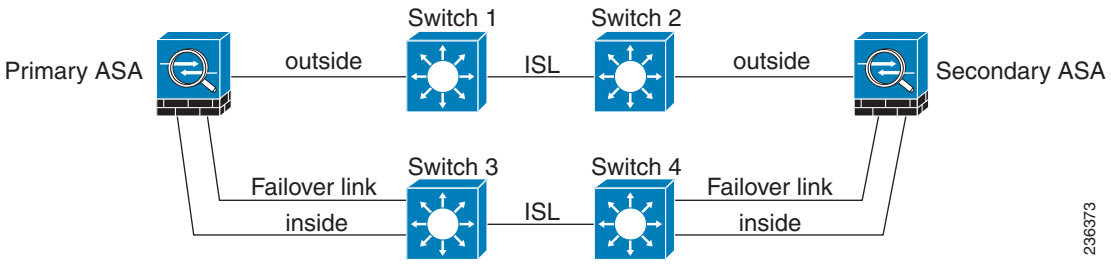
Figure 9-4 Connecting with a Cable



Scenario 3—Recommended

If the ASA data interfaces are connected to more than one set of switches, then a failover link can be connected to one of the switches, preferably the switch on the secure (inside) side of network, as shown in Figure 9-5.

Figure 9-5 Connecting with a Secure Switch



Scenario 4—Recommended

The most reliable failover configurations use a redundant interface on the failover link, as shown in Figure 9-6 and Figure 9-7.

Figure 9-6 Connecting with Redundant Interfaces

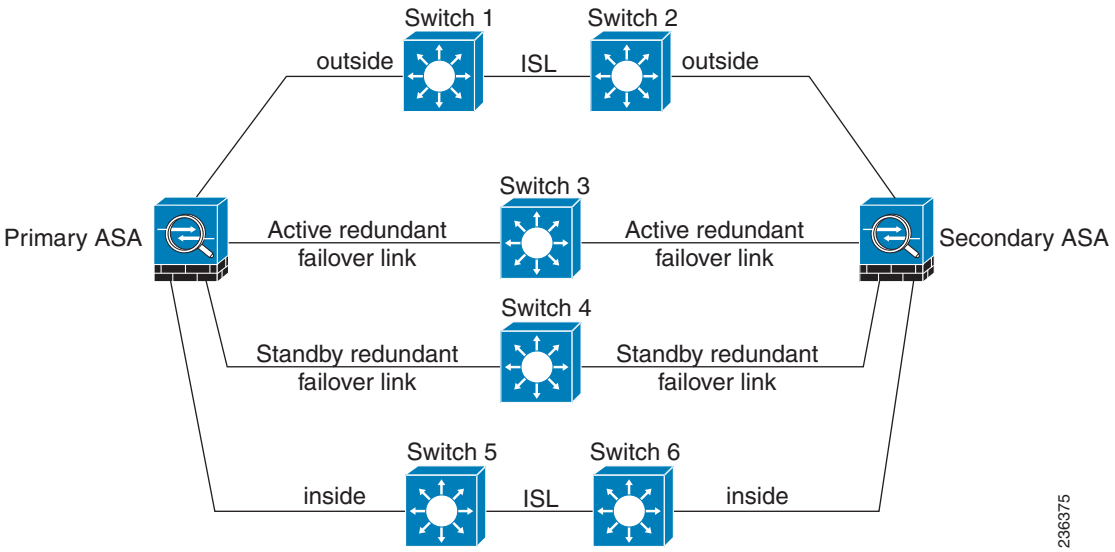
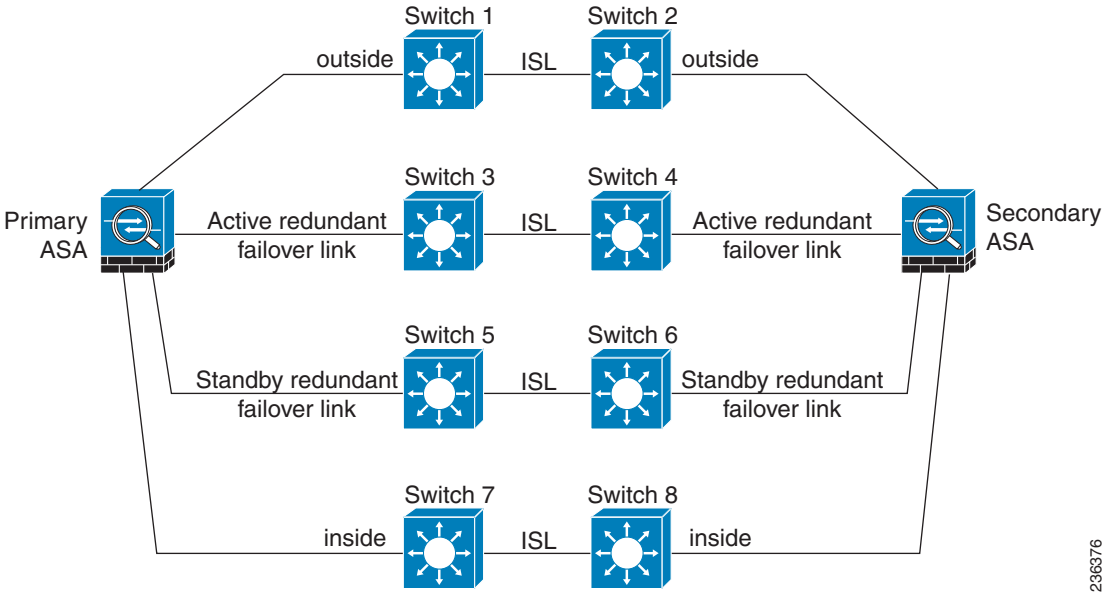


Figure 9-7 Connecting with Inter-switch Links



MAC Addresses and IP Addresses

When you configure your interfaces, you must specify an active IP address and a standby IP address on the same network.

1. When the primary unit or failover group fails over, the secondary unit assumes the IP addresses and MAC addresses of the primary unit and begins passing traffic.
2. 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 If the secondary unit boots without detecting the primary unit, the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. However, when the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic.Similarly, if you swap out the primary unit with new hardware, a new MAC address is used.

Virtual MAC addresses guard against this disruption because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. In multiple context mode, the ASA generates virtual active and standby MAC addresses by default. See the "Information About MAC Addresses" section on page 8-11 for more information. In single context mode, you can manually configure virtual MAC addresses; see the "Configuring Active/Active Failover" section on page 9-32 for more information.

If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The ASA does not send gratuitous ARPs for static NAT addresses when the MAC address changes, so connected routers do not learn of the MAC address change for these addresses.

Note The IP address and MAC address for the state link do not change at failover; the only exception is if the state link is configured on a regular data interface.

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 9-8
- Inter-Chassis Failover, page 9-9

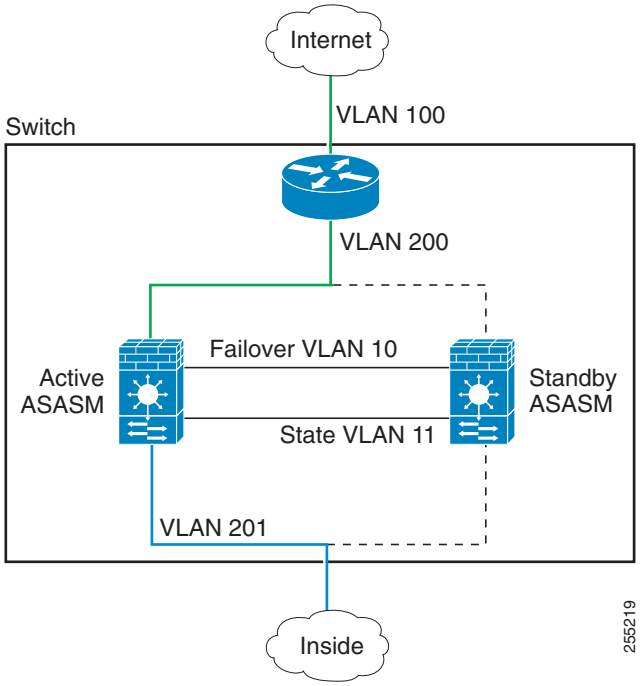
Intra-Chassis Failover

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 9-9.

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 9-8 shows a typical intra-switch configuration.

Figure 9-8 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.

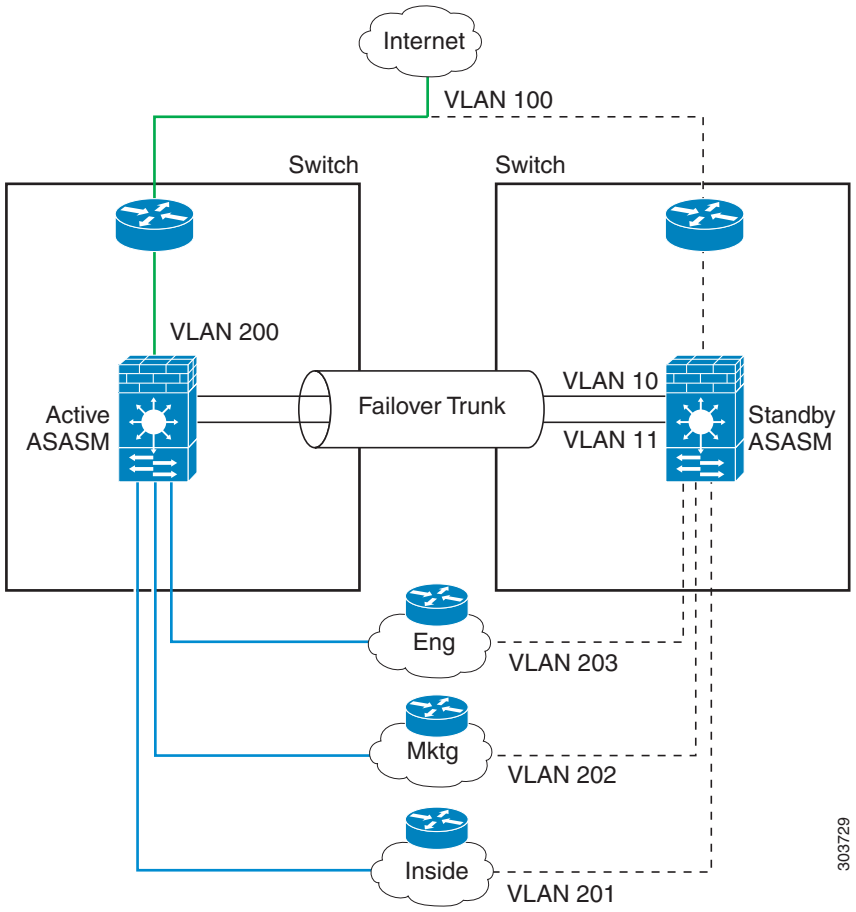
For the best reliability of failover communications between ASASMs, we recommend that you configure an EtherChannel trunk port between the two switches to carry the failover and state VLANs.

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 9-9 shows a typical switch and ASASM redundancy configuration. The trunk between the two switches carries the failover ASASM VLANs (VLANs 10 and 11).

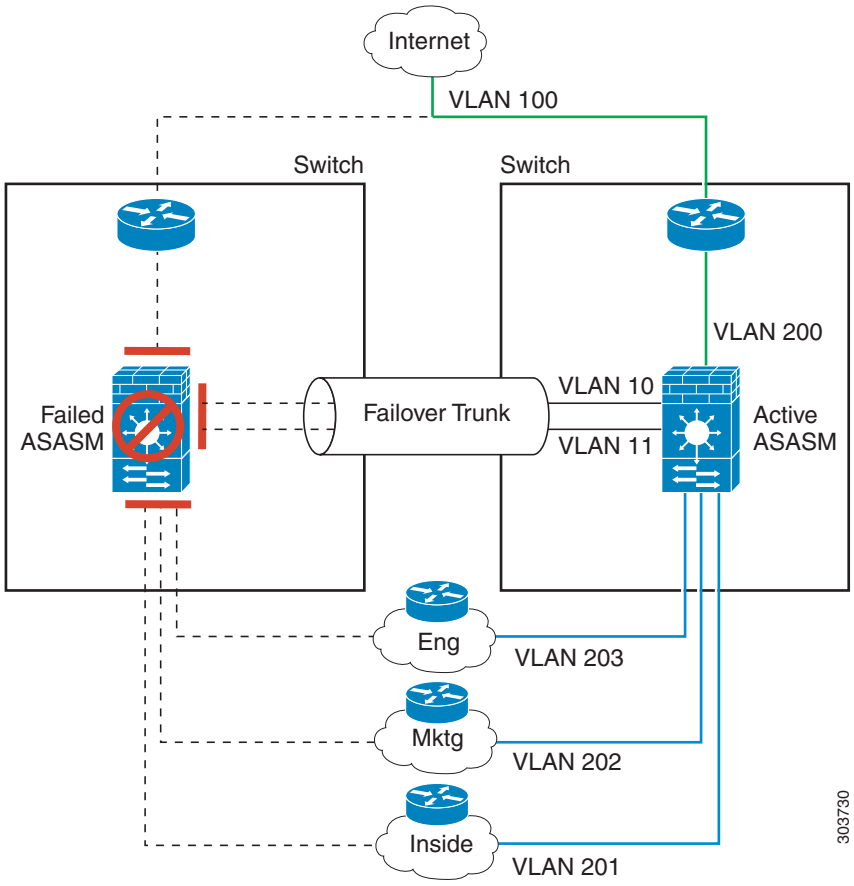
Note ASASM failover is independent of the switch failover operation; however, ASASM works in any switch failover scenario.

Figure 9-9 Normal Operation



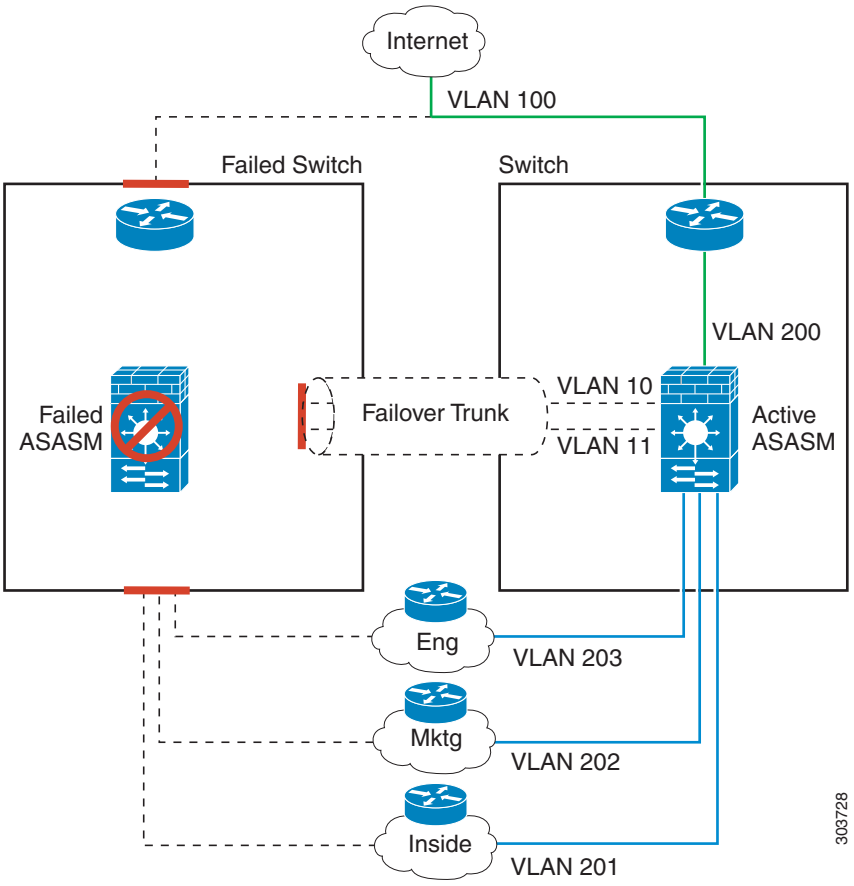
If the primary ASASM fails, then the secondary ASASM becomes active and successfully passes the firewall VLANs (Figure 9-10).

Figure 9-10 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 9-11).

Figure 9-11 Switch Failure



Stateless and Stateful Failover

The ASA supports two types of failover, stateless and stateful for both the Active/Standby and Active/Active modes.

- Stateless Failover, page 9-13
- Stateful Failover, page 9-13

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 failover is not recommended for clientless SSL VPN.

Stateless 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, or in Active/Active failover, between the active and standby failover groups. 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.

- Supported Features, page 9-13
- Unsupported Features, page 9-14

Supported Features

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)—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 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.
- Dynamic Routing Protocols—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.

- Cisco IP SoftPhone sessions—If a 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 Call Manager. This connection loss 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 Call Manager unreachable and unregisters itself.
- VPN—VPN end-users do not have to reauthenticate or reconnect the VPN session after a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.

Unsupported Features

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
- Application inspections that are subject to advanced TCP-state tracking—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, such as the ASA IPS SSP or ASA CX SSP.
- Phone proxy connections—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.
- Selected clientless SSL VPN features:
  - Smart Tunnels
  - Port Forwarding
  - Plugins
  - Java Applets
  - IPv6 clientless or Anyconnect sessions
  - Citrix authentication (Citrix users must reauthenticate after failover)

Transparent Firewall Mode Requirements

- Transparent Mode Requirements for Appliances, page 9-14
- Transparent Mode Requirements for Modules, page 9-15

Transparent Mode Requirements for Appliances

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 with an EtherType access rule.

  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:

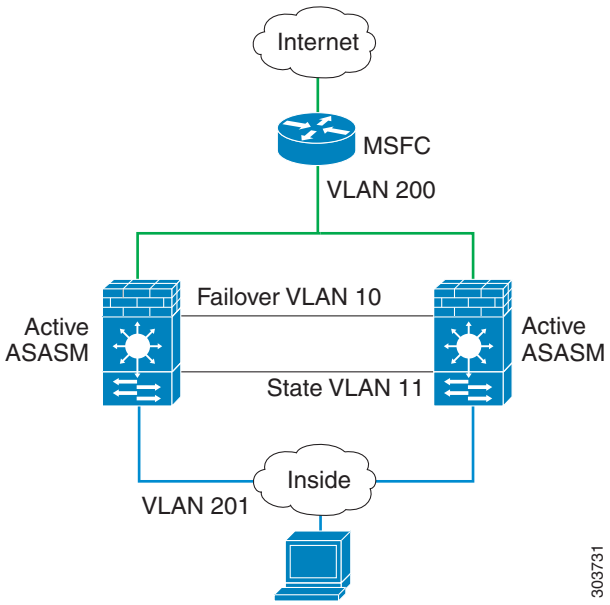
- Disable interface monitoring.
- Increase interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the interface holdtime.

Transparent Mode Requirements for Modules

To avoid loops when you use failover in transparent mode, you should allow BPDUs to pass (the default), and you must use switch software that supports BPDU forwarding.

Loops can occur if both modules are active at the same time, such as when both modules are discovering each other's presence, or due to a bad failover link. Because the ASASMs bridge packets between the same two VLANs, loops can occur when inside packets destined for the outside get endlessly replicated by both ASASMs (see Figure 9-12). The spanning tree protocol can break such loops if there is a timely exchange of BPDUs. To break the loop, BPDUs sent between VLAN 200 and VLAN 201 need to be bridged.

Figure 9-12 Transparent Mode Loop



Failover Health Monitoring

The ASA monitors each unit for overall health and for interface health. This section includes information about how the ASA performs tests to determine the state of each unit.

- Unit Health Monitoring, page 9-16
- Interface Monitoring, page 9-16

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 data interface, including the failover link, to validate whether or not the peer is responsive. The action that the ASA takes depends on the response from the other unit. See the following possible actions:

- If the ASA receives a response on the failover link, 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 a data 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.

Interface Monitoring

You can monitor up to 250 interfaces (in multiple mode, divided between all contexts). You should monitor important interfaces. For example in multiple mode, 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.

Monitored 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.

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 9-1 shows the minimum, default, and maximum failover times.

Table 9-1 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

Configuration Synchronization

Failover includes two types of configuration synchronization:

- Running Configuration Replication, page 9-18
- Command Replication, page 9-18

Running Configuration Replication

Running configuration replication 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.

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." Depending on the size of the configuration, replication can take from a few seconds to several minutes.

On the standby unit, the configuration exists only in running memory. You should save the configuration to flash memory.

Note 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 during the configuration replication process.

Note The **crypto ca server** command and related sub commands are not synchronized to the failover peer.

Note Configuration syncing does not replicate the following files and configuration components, so you must copy these files manually so they match:

- AnyConnect images
- CSD images
- AnyConnect profiles
- Local Certificate Authorities (CAs)
- ASA images
- ASDM images

Command Replication

After startup, commands that you enter on the active unit are immediately replicated to the standby unit.You do not have to save the active configuration to flash memory to replicate the commands.

In Active/Active failover,changes entered in the system execution space are replicated from the unit on which failover group 1 is in the active 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.

The following commands 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 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**

Information About Active/Standby Failover

Active/Standby failover lets you 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.

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 Roles and Active/Standby Status, page 9-19
- Active Unit Determination at Startup, page 9-19
- Failover Events, page 9-20

Primary/Secondary Roles 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.

Active Unit Determination at Startup

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.

Failover Events

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 9-2 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 9-2 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 link as failed	Mark failover link 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 link as failed	Become active	If the failover link is down at startup, both units become active.
State link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
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.

Information About Active/Active Failover

This section describes Active/Active failover. This section includes the following topics:

- Active/Active Failover Overview, page 9-21

- Primary/Secondary Roles and Active/Standby Status for a Failover Group, page 9-21
- Failover Events, page 9-22

Active/Active Failover Overview

In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into a maximum of 2 failover groups.

A failover group is simply a logical group of one or more security contexts. You can assign failover group to be active on the primary ASA, and failover group 2 to be active on the secondary ASA. When a failover occurs, it occurs at the failover group level. For example, depending on interface failure patterns, it is possible for failover group 1 to fail over to the secondary ASA, and subsequently failover group 2 to fail over to the primary ASA. This event could occur if the interfaces in failover group 1 are down on the primary ASA but up on the secondary ASA, while the interfaces in failover group 2 are down on the secondary ASA but up on the primary ASA.

The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default. If you want Active/Active failover, but are otherwise uninterested in multiple contexts, the simplest configuration would be to add one additional context and assign it to failover group 2.

Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.

Note You can assign both failover groups to one ASA if desired, but then you are not taking advantage of having two active ASAs.

Primary/Secondary Roles and Active/Standby Status for a Failover Group

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:

- The primary unit provides the running configuration to the pair when they boot simultaneously.
- Each failover group in the configuration is configured with a primary or secondary unit preference.

Active Unit Determination for Failover Groups at Startup

The unit on which a failover group becomes active 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.

Failover Events

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.

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.

Table 9-3 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 9-3 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.
State 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 link 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.

Licensing Requirements Failover

Active/Standby Failover

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.

Active/Active Failover

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 Failover

See the "Failover System Requirements" section on page 9-2.

Guidelines and Limitations

For Auto Update guidelines with failover, see the "Auto Update Server Support in Failover Configurations" section on page 46-32.

Context Mode Guidelines

- Active/Standby mode is supported in single and multiple context mode.
- Active/Active mode is supported only in multiple context mode.
- For multiple context mode, perform all steps in the system execution space unless otherwise noted.
- 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 in each context sequentially.

Firewall Mode Guidelines

Supported in transparent and routed firewall mode.

IPv6 Guidelines

IPv6 is supported.

Model Guidelines

Stateful failover is not supported on the ASA 5505. See the "Licensing Requirements Failover" section on page 9-23 for other guidelines.

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 problem occurs when 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.
- You can monitor up to 250 interfaces on a unit, across all contexts.
- For Active/Active failover, no two interfaces in the same context should be configured in the same ASR group.
- For Active/Active failover, you can define a maximum of two failover groups.
- For Active/Active failover, 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.

Default Settings

By default, the failover policy consists of the following:

- No HTTP replication in Stateful Failover.
- A single interface failure causes failover.
- The interface poll time is 5 seconds.
- The interface hold time is 25 seconds.
- The unit poll time is 1 second.

- The unit hold time is 15 seconds.
- Virtual MAC addresses are enabled in multiple context mode; in single context mode, they are disabled.
- Monitoring on all physical interfaces, or for the ASA 5505 and ASASM, all VLAN interfaces.

Configuring Active/Standby Failover

The High Availability and Scalability Wizard guides you through a step-by-step process of creating an Active/Standby failover configuration.

- Detailed Steps 1—Starting the Wizard, page 9-25
- Detailed Steps 2—Failover Peer Connectivity and Compatibility Check, page 9-26
- Detailed Steps 3—LAN Link Configuration, page 9-27
- Detailed Steps 4—State Link Configuration, page 9-29
- Detailed Steps 5—Standby Address Configuration, page 9-29
- Detailed Steps 6—Summary, page 9-30

Detailed Steps 1—Starting the Wizard

Step 1 Choose **Wizards > High Availability and Scalability**.



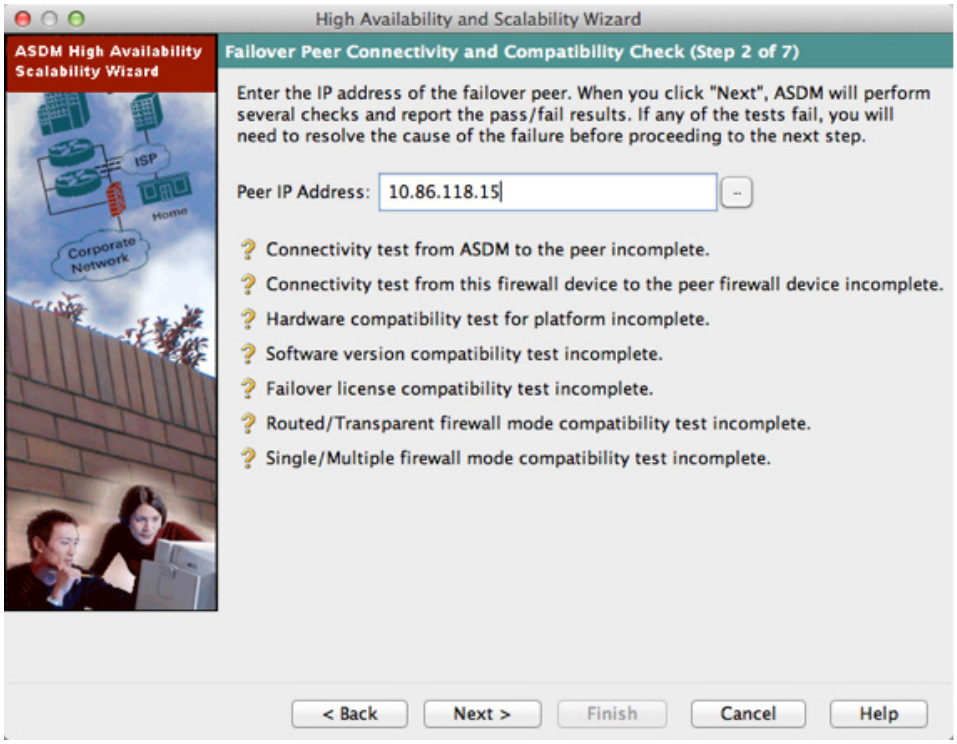
Step 2 In the Configuration Type screen, click **Configure Active/Standby failover**, and click **Next**.

The Failover Peer Connectivity and Compatibility screen appears. See "Detailed Steps 2—Failover Peer Connectivity and Compatibility Check".

Detailed Steps 2—Failover Peer Connectivity and Compatibility Check

Step 1 In the Peer IP Address field, enter the IP address of the peer unit. This address must be an interface that has ASDM access enabled on it.

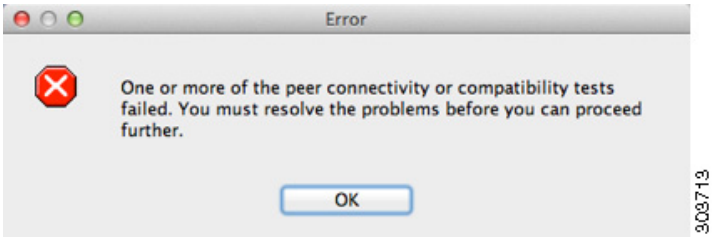
By default, the peer address is assigned to be the standby address for the ASDM management interface.



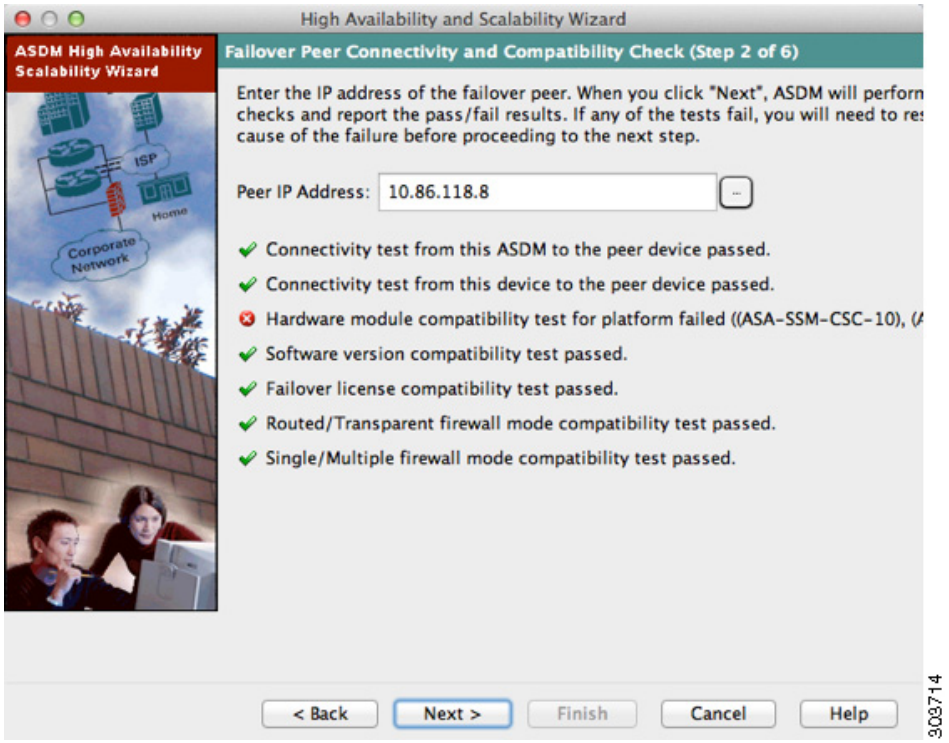
Step 2 Click **Next** to perform connectivity and compatibility tests. You are prompted to log into the peer unit.

If the tests succeed, the LAN Link Configuration screen appears. See "Detailed Steps 3—LAN Link Configuration".

If any of the tests fail, you see an error dialog box.

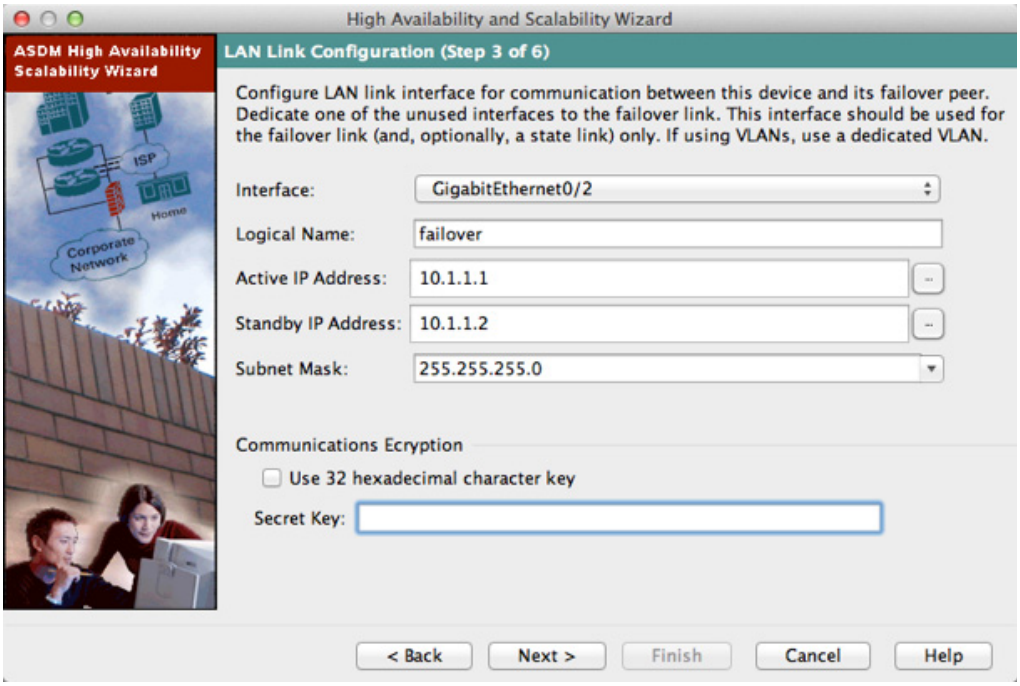


After you click OK, you are returned to the compatibility check screen, which shows which tests failed. Click **Cancel** to exit the wizard and resolve any issues before trying again.



Detailed Steps 3—LAN Link Configuration

Step 1 Configure the failover link parameters:



a. Interface—Choose the interface to use for failover communication.

b. Logical Name—Enter a name for the interface.

c. Active IP Address—Enter the IP address used for the failover link on the primary unit. This should be on an unused subnet.

d. Standby IP Address—Enter the IP address used for the failover link on the secondary unit, on the same network as the active IP address.

e. Subnet Mask—Enter or choose a subnet mask for the Active IP and Standby IP addresses.

f. (ASA 5505 only) Switch Port—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. By default, VLAN 1 is the inside interface, so you should choose a different VLAN.

Note To provide sufficient bandwidth for failover, do not use trunks or PoE for failover.

g. (Optional) Communications Encryption—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see the "Modifying the Failover Setup" section on page 9-47).

- Secret Key—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.

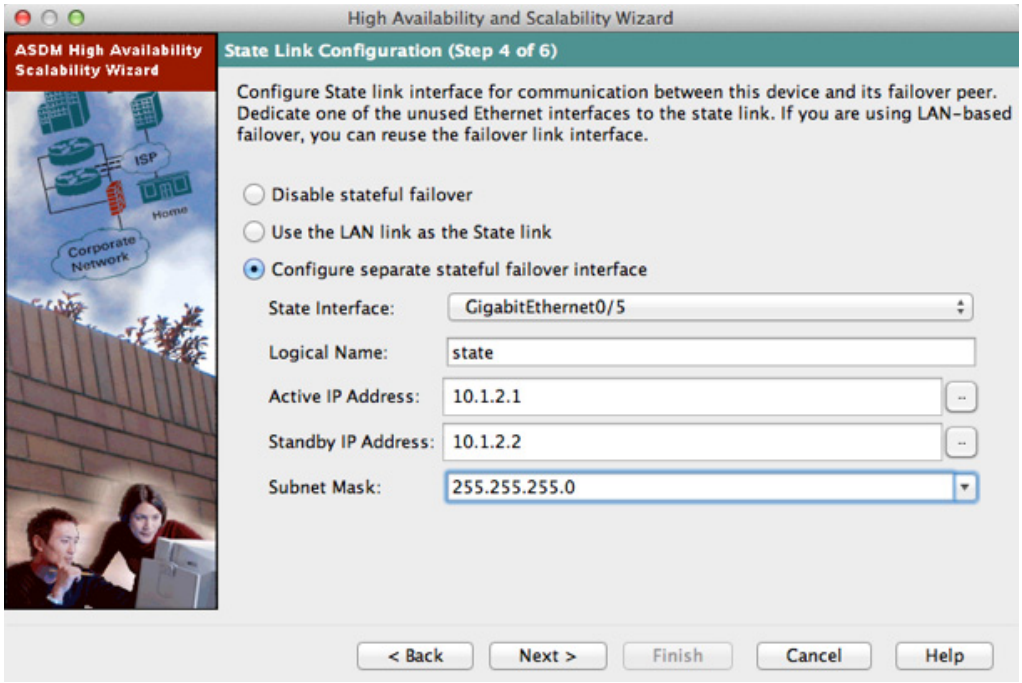
  Use 32 hexadecimal character key—To use a 32-hexadecimal key for the secret key, check this check box.

Step 2 Click **Next**.

The State Link Configuration screen appears. See "Detailed Steps 4—State Link Configuration".

Detailed Steps 4—State Link Configuration

Step 1 Choose one of the following options for the state link:



- **Disable Stateful Failover**—Disables Stateful Failover.
- **Use the LAN link as the State Link**—Passes state information across the failover link.
- **Configure another interface for Stateful failover**—Configures an unused interface as the state link.

Step 2 If you choose another interface for Stateful Failover, configure the following parameters:

a. State interface—Choose an unused interface.

b. Logical Name—Enter the name for the state link.

c. Active IP Address—Enter the IP address for the state link on the primary unit. This should be on an unused subnet, different from the failover link.

d. Standby IP Address—Enter the IP address for the state link on the secondary unit, on the same network as the active IP address.

e. Subnet Mask—Enter or choose a subnet mask for the Active IP and Standby IP addresses.

Step 3 Click **Next**.

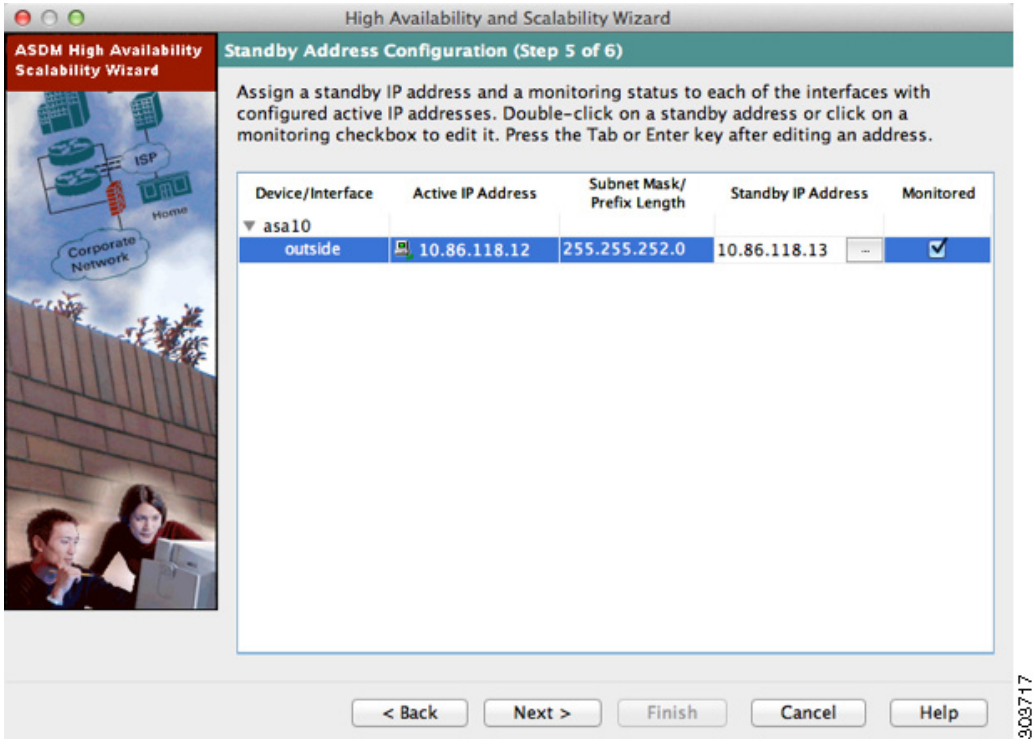
The Standby Address Configuration screen appears. See "Detailed Steps 5—Standby Address Configuration".

Detailed Steps 5—Standby Address Configuration

Step 1 Assign standby IP addresses to the data interfaces on the ASA. Any currently configured interfaces appear.

By default, the peer address that you specified on the Failover Peer Connectivity and Compatibility screen is assigned to be the standby address for the ASDM management interface.

If you configure data interfaces later, you can assign standby IP addresses at that time, or on the Configuration > Device Management > High Availability > Failover > Interfaces tab (see the "Configuring Interface Monitoring and Standby Addresses" section on page 9-44).



a. Select the **Active IP Address** field to edit or add an active IP address.

b. Select the **Standby IP Address** field to edit or add a standby IP address.

c. Select the **Subnet Mask/Prefix Length** field to edit the subnet mask or prefix length.

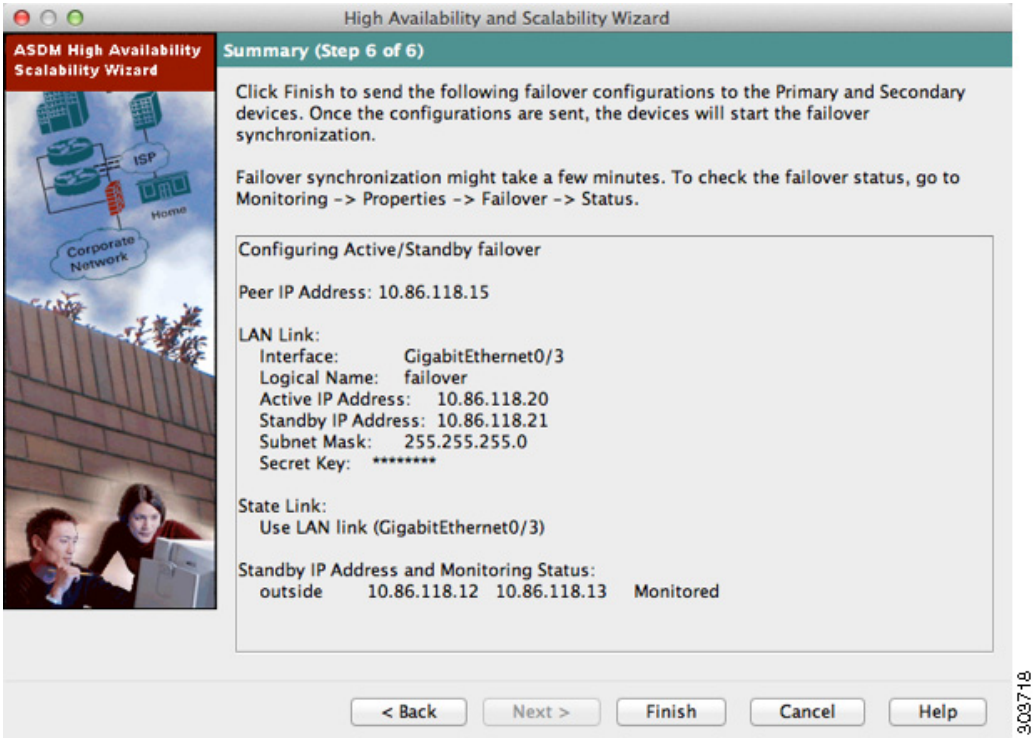
d. Check the **Monitored** check box to enable health monitoring for that interface. Uncheck the check box to disable health monitoring. By default, health monitoring of physical interfaces is enabled, and health monitoring of subinterfaces is disabled.

Step 2 Click **Next**.

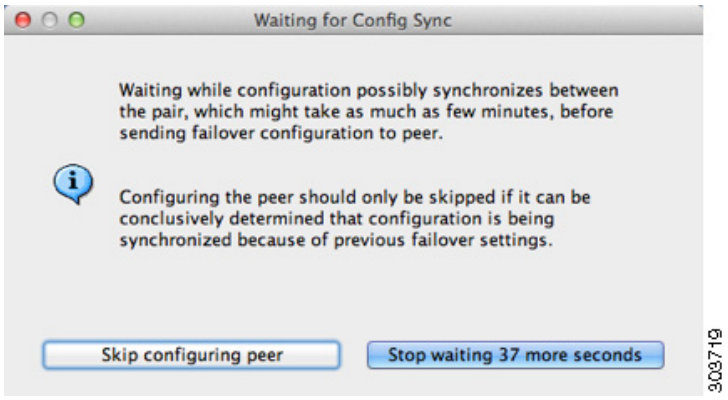
The Summary screen appears. See "Detailed Steps 6—Summary".

Detailed Steps 6—Summary

Step 1 Verify your settings and click **Finish** to send your configuration to the primary unit.



The wizard shows the Waiting for Config Sync screen.



After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting *xx* more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.

Step 2 Click **OK**.

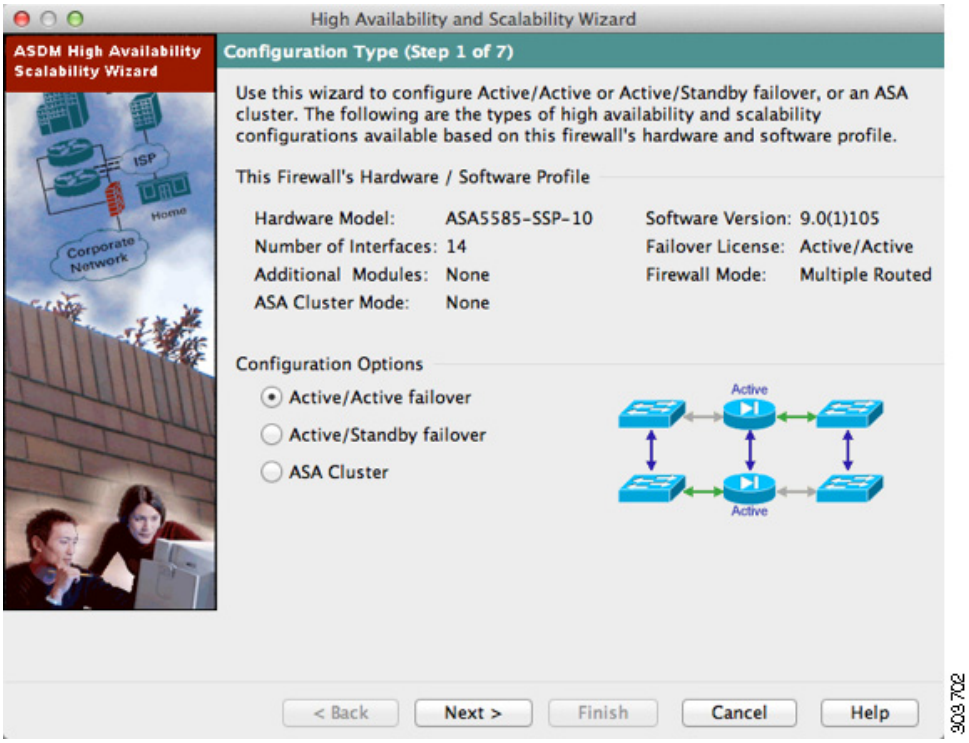
Configuring Active/Active Failover

The High Availability and Scalability Wizard guides you through a step-by-step process of creating an Active/Active failover configuration.

- Detailed Steps 1—Starting the Wizard, page 9-32
- Detailed Steps 2—Failover Peer Connectivity and Compatibility Check, page 9-33
- Detailed Steps 3—Security Context Configuration, page 9-35
- Detailed Steps 4—LAN Link Configuration, page 9-36
- Detailed Steps 5—State Link Configuration, page 9-37
- Detailed Steps 6—Standby Address Configuration, page 9-38
- Detailed Steps 7—Summary, page 9-39

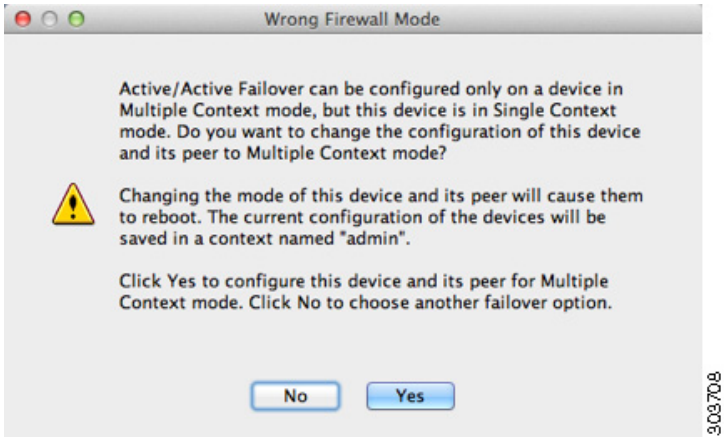
Detailed Steps 1—Starting the Wizard

Step 1 Choose **Wizards > High Availability and Scalability**.



Step 2 In the Configuration Type screen, click **Configure Active/Active failover**, and click **Next**.

- If your devices are already in multiple context mode, the Failover Peer Connectivity and Compatibility screen appears.
- If your devices are not yet in multiple context mode, you see the Wrong Firewall Mode dialog box. Click **Yes** to change the mode as part of the wizard, or click **No** to exit the wizard. If you click Yes, you are returned to the Configuration Type screen. Click **Next**. the Failover Peer Connectivity and Compatibility screen appears. For more information about multiple context mode, see Chapter 8, "Configuring Multiple Context Mode."

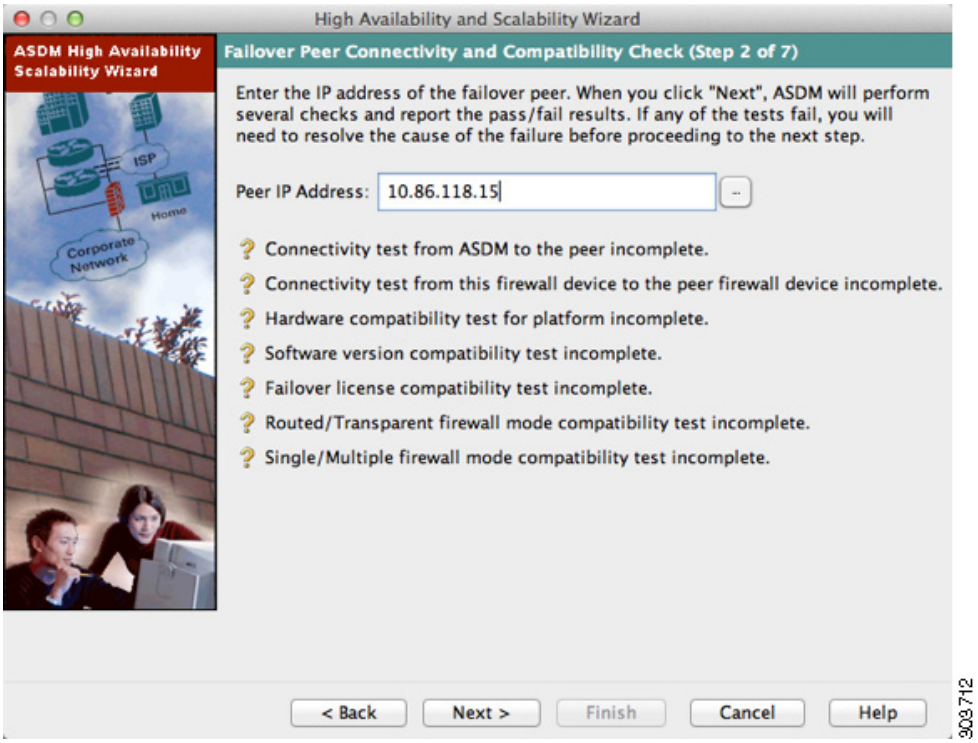


See "Detailed Steps 2—Failover Peer Connectivity and Compatibility Check".

Detailed Steps 2—Failover Peer Connectivity and Compatibility Check

Step 1 In the Peer IP Address field, enter the IP address of the peer unit. This address must be an interface that has ASDM access enabled on it.

By default, the peer address is assigned to be the standby address for the interface to which ASDM is connected.



Step 2 Click **Next** to perform the following connectivity and compatibility tests:

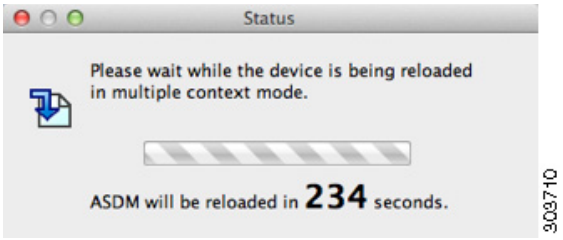
- Connectivity test from this ASDM to the peer unit
- Connectivity test from this firewall device to the peer firewall device
- Hardware compatibility test for the platform
- Software version compatibility
- Failover license compatibility
- Firewall mode compatibility (routed or transparent)
- Context mode compatibility (single or multiple)

Step 3 You are prompted to log into the peer unit.

- If you opted to change to multiple context mode, you see the Wrong Firewall Mode dialog box.

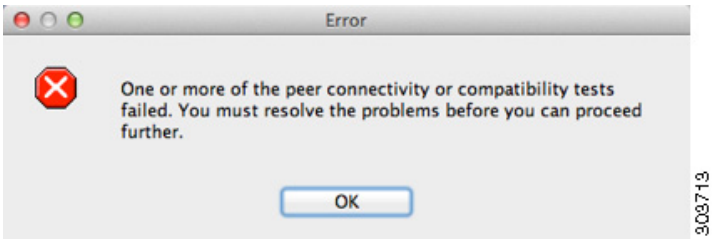


Click **Yes** to change the mode and reload both units. You see the Status dialog box showing a countdown while ASDM waits for the units to reload.

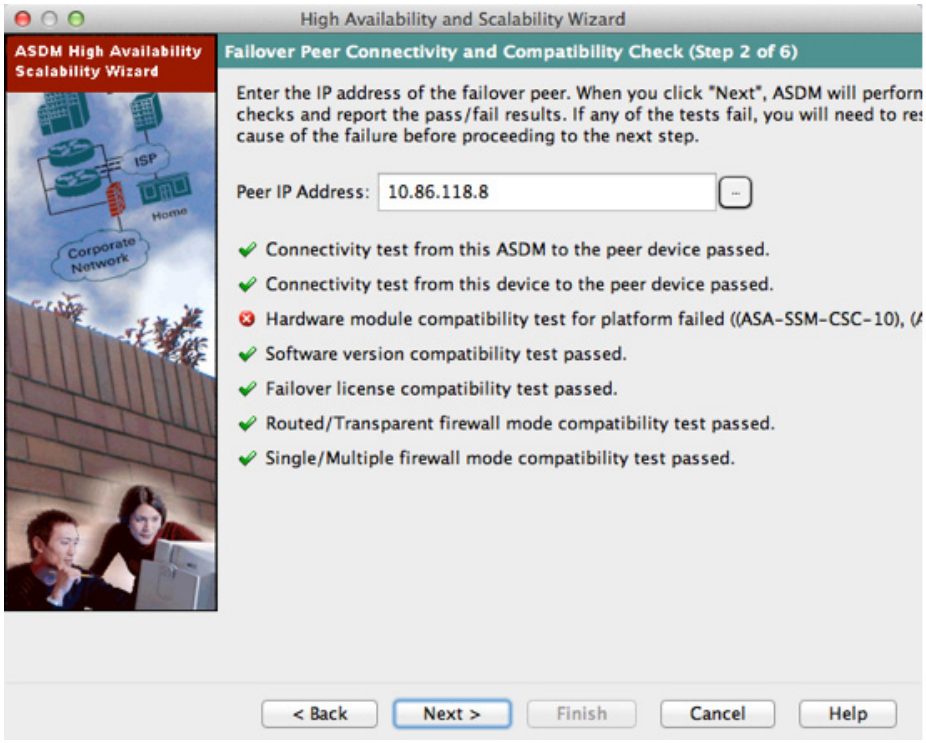


At the end of the countdown, ASDM reconnects to the primary unit and you return to the Failover Peer Connectivity and Compatibility screen. Click **Next** to recheck compatibility.

- If the tests succeed, the Security Context Configuration screen appears. See "Detailed Steps 3—Security Context Configuration".
- If any of the tests fail, you see an error dialog box.

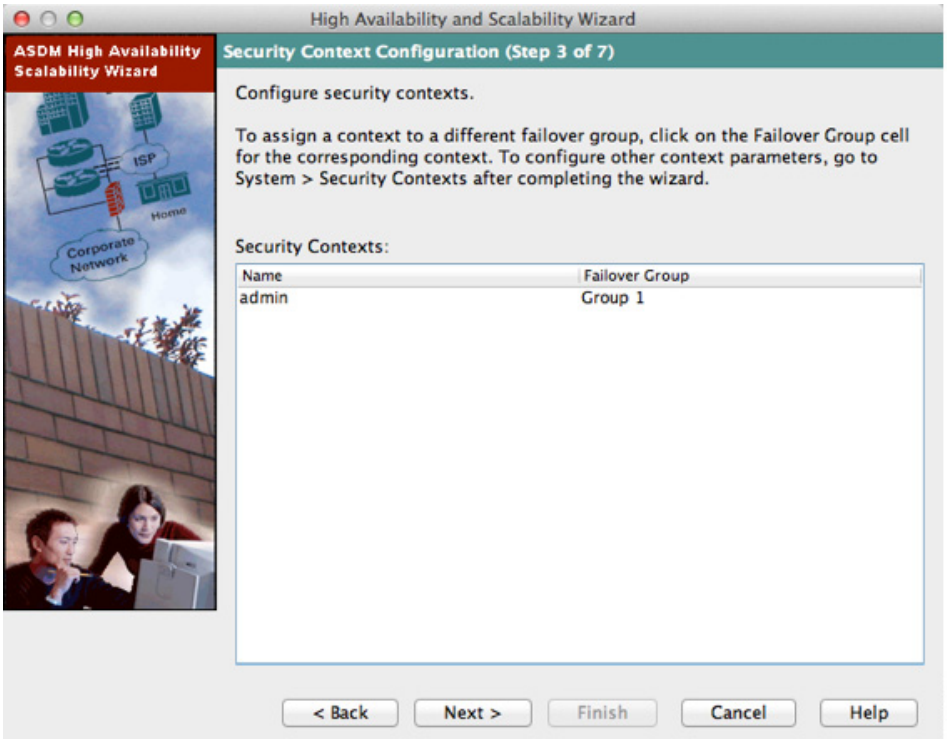


After you click OK, you are returned to the compatibility check screen, which shows which tests failed. Click **Cancel** to exit the wizard and resolve any issues before trying again.



Detailed Steps 3—Security Context Configuration

Step 1 For existing contexts, you can set the failover group (1 or 2). If you converted to multiple context mode as part of the wizard, you will only see the admin context. You can add other contexts after you exit the wizard.

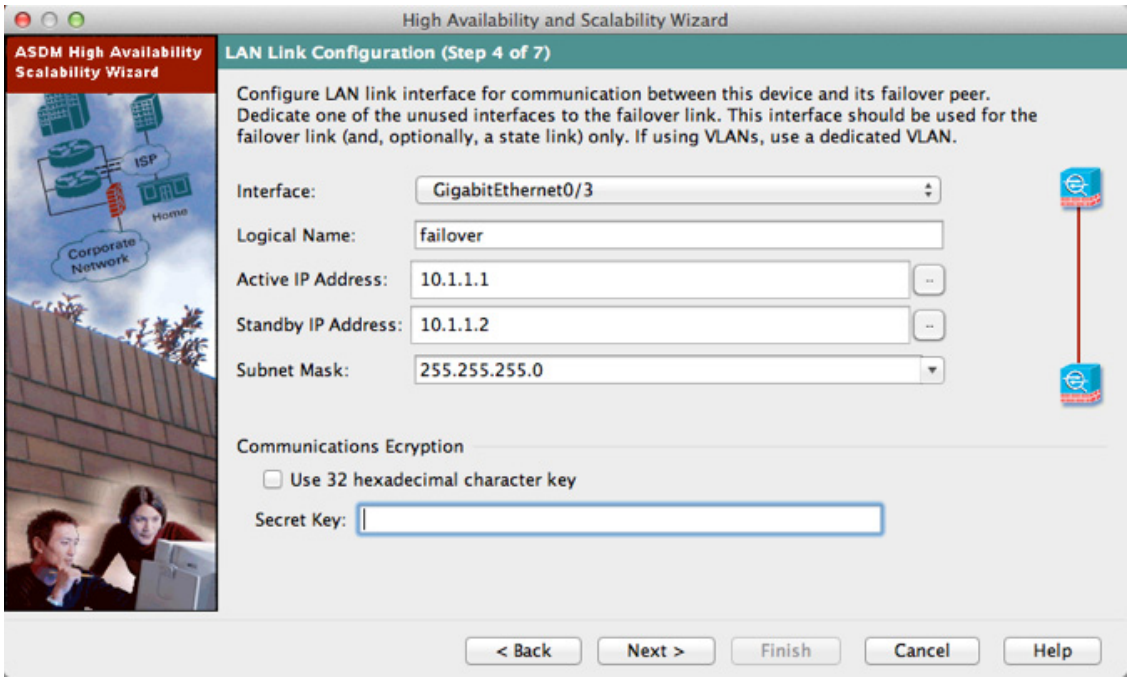


Step 2 Click **Next**.

The LAN Link Configuration screen appears. See "Detailed Steps 4—LAN Link Configuration".

Detailed Steps 4—LAN Link Configuration

Step 1 Configure the failover link parameters:



a. Interface—Choose the interface to use for failover communication.

b. Logical Name—Enter a name for the interface.

c. Active IP Address—Enter the IP address used for the failover link on the primary unit. This IP address should be on an unused subnet.

d. Standby IP Address—Enter the IP address used for the failover link on the secondary unit.

e. Subnet Mask—Enter or choose a subnet mask for the Active IP and Standby IP addresses.

f. (Optional) Communications Encryption—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see the "Modifying the Failover Setup" section on page 9-47).

- Secret Key—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.

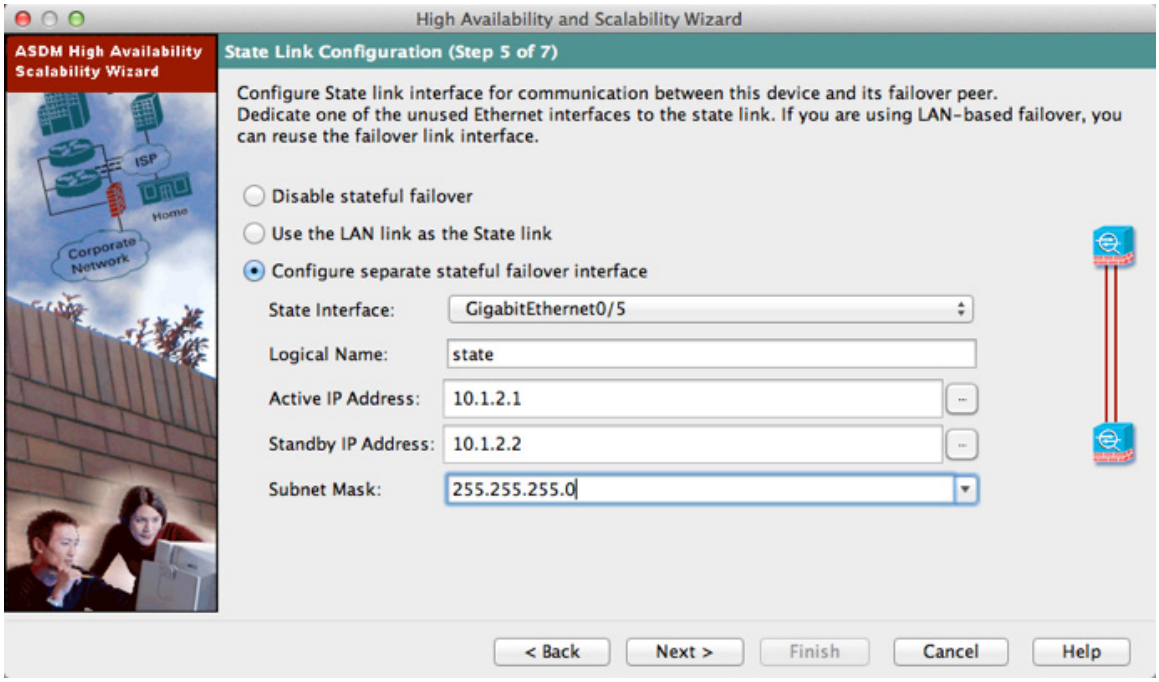
  Use 32 hexadecimal character key—To use a 32-hexadecimal key for the secret key, check this check box.

Step 2 Click **Next**.

The State Link Configuration screen appears. See "Detailed Steps 5—State Link Configuration".

Detailed Steps 5—State Link Configuration

Step 1 Choose one of the following options for the state link:



- **Disable Stateful Failover**—Disables Stateful Failover.
- **Use the LAN link as the State Link**—Passes state information across the failover link.
- **Configure another interface for Stateful failover**—Configures an unused interface as the state link.

Step 2 If you choose another interface for Stateful Failover, configure the following parameters:

a. State interface—Choose an unused interface.

b. Logical Name—Enter the name for the state link. For example, change the name to "state."

c. Active IP Address—Enter the IP address for the state link on the primary unit. This should be on an unused subnet, different from the failover link.

d. Standby IP Address—Enter the IP address for the state link on the secondary unit.

e. Subnet Mask—Enter or choose a subnet mask for the Active IP and Standby IP addresses.

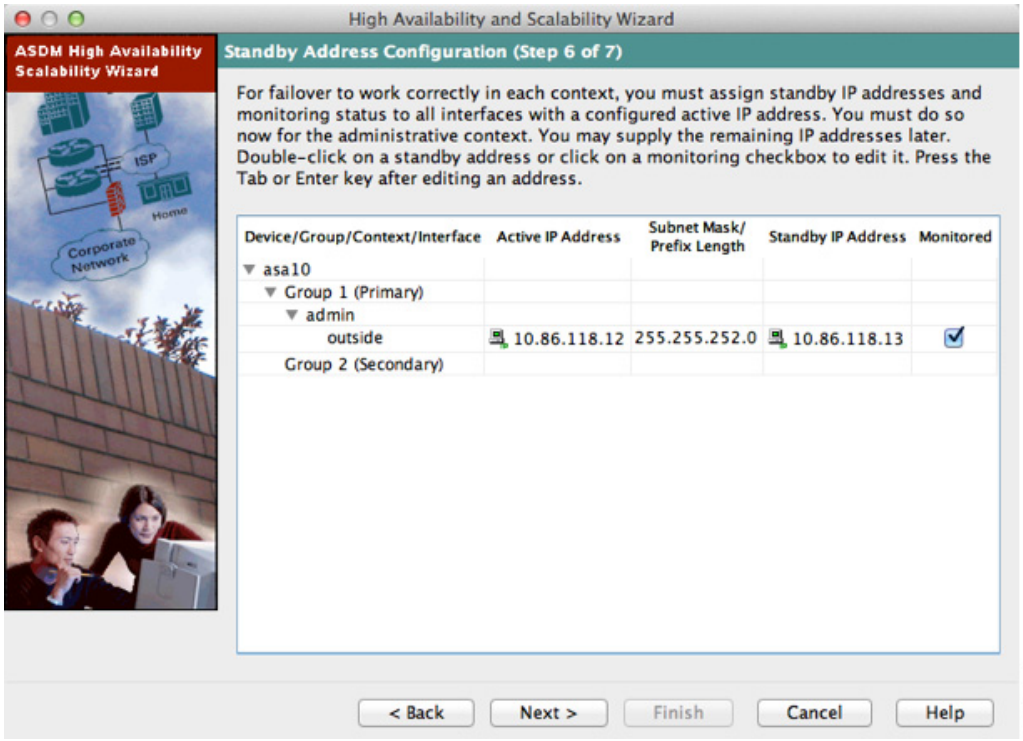
Step 3 Click **Next**.

The Standby Address Configuration screen appears. See "Detailed Steps 6—Standby Address Configuration".

Detailed Steps 6—Standby Address Configuration

Step 1 Assign standby IP addresses to the interfaces on the ASA. The interfaces currently configured on the failover devices appear.

By default, the peer address that you specified on the Failover Peer Connectivity and Compatibility screen is assigned to be the standby address for the interface to which ASDM is connected.



a. Select the **Active IP Address** 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.

b. Select the **Standby IP Address** 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.

c. Select the **Subnet Mask/Prefix Length** field to edit the subnet mask or prefix length.

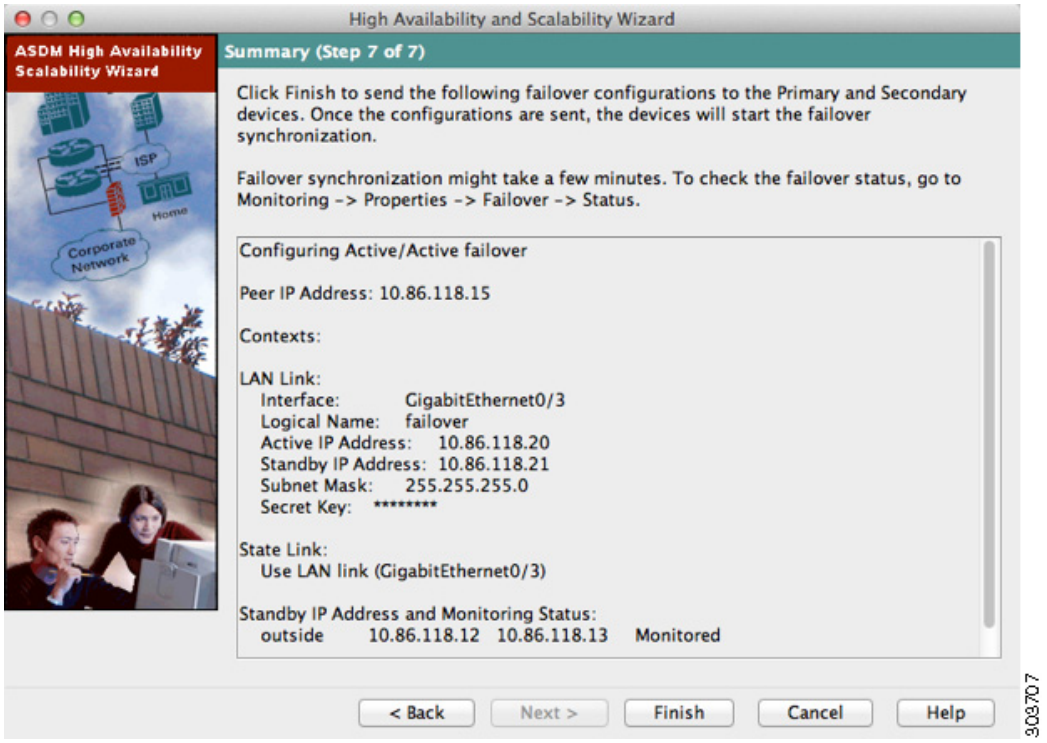
d. Check the **Monitored** check box to enable health monitoring for that interface. Uncheck the check box to disable health monitoring. By default, health monitoring of physical interfaces is enabled, and health monitoring of subinterfaces is disabled.

Step 2 Click **Next**.

The Summary screen appears. See "Detailed Steps 7—Summary".

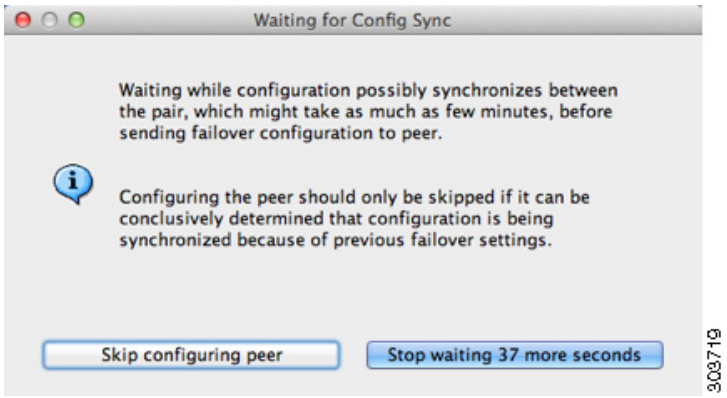
Detailed Steps 7—Summary

Step 1 Verify your settings and click **Finish** to send your configuration to the primary device.



The wizard shows the Waiting for Config Sync screen. After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting *xx* more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.



Step 2 Click **OK**.

Configuring Optional Failover Parameters

You can customize failover settings as desired.

- Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses, page 9-41
- Configuring Interface Monitoring and Standby Addresses, page 9-44
- Configuring Support for Asymmetrically Routed Packets (Active/Active Mode), page 9-45

Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses

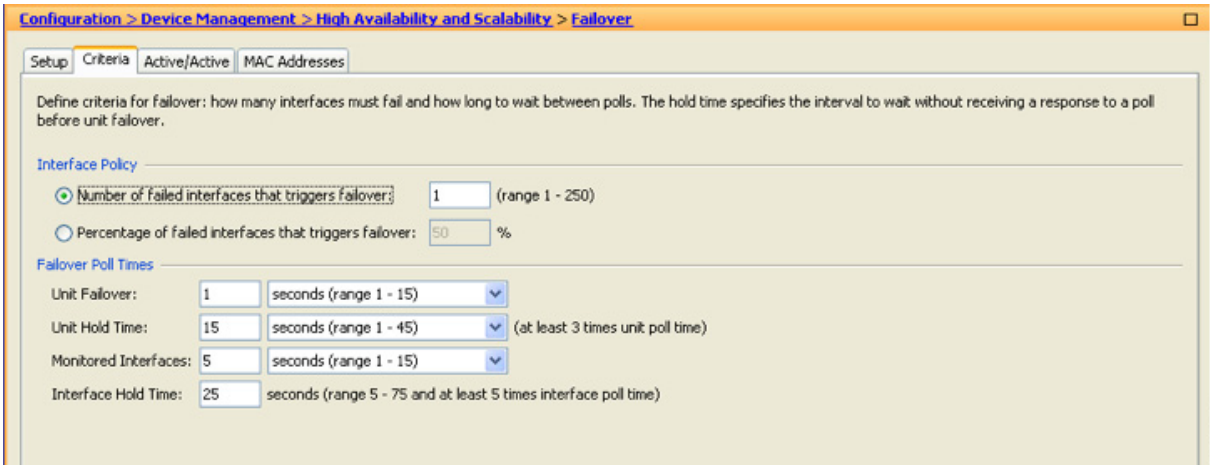
See the "Default Settings" section on page 9-24 for the default settings for many parameters that you can change in this section. For Active/Active mode, you set most criteria per failover group. This section includes enabling HTTP replication per failover group for Active/Active mode; to configure HTTP replication for Active/Standby mode, see the "Modifying the Failover Setup" section on page 9-47.

Prerequisites

Configure these settings in the system execution space in multiple context mode.

Detailed Steps

Step 1 Choose **Configuration > Device Management > High Availability > Failover > Criteria**.

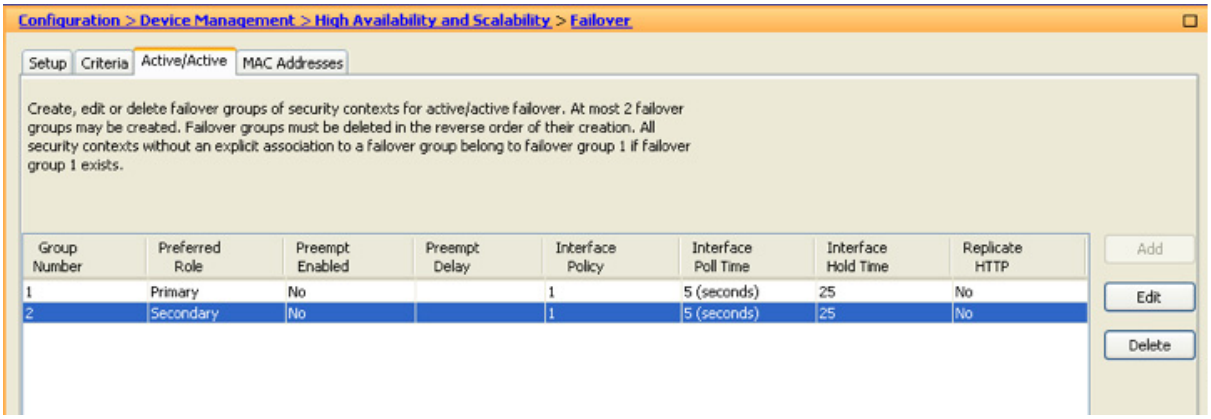


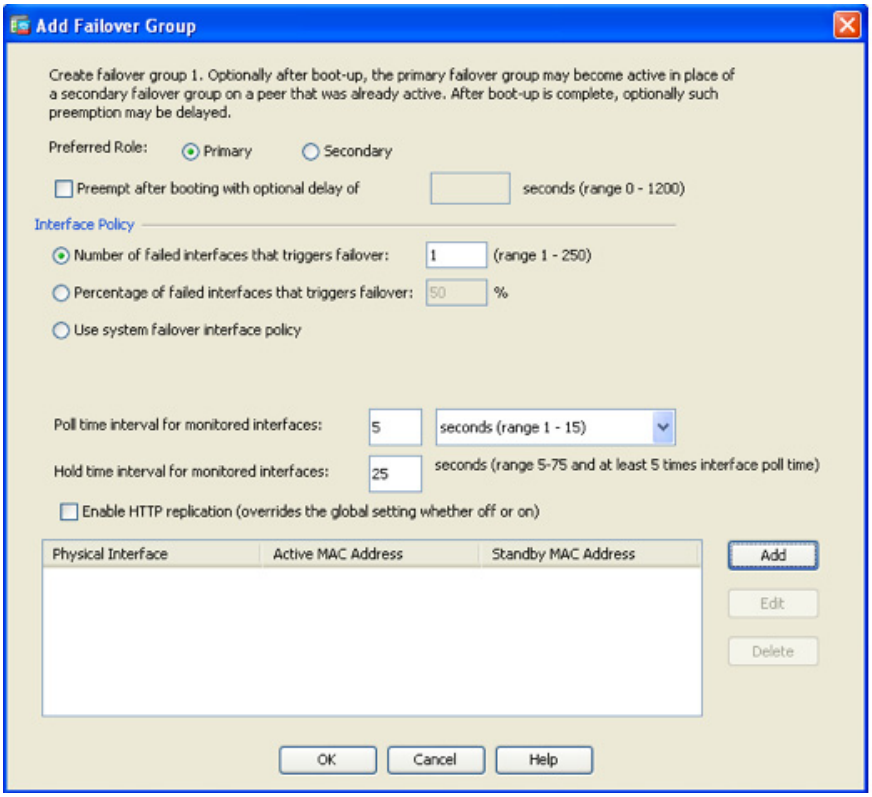
Step 2 In the Failover Poll Times area, configure the unit poll times:

- 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.

Note Other settings on this pane apply only to Active/Standby mode. In Active/Active mode, you must configure the rest of the parameters per failover group.

Step 3 (Active/Active mode only) Choose **Configuration > Device Management > High Availability > Failover > Active/Active**, then choose a failover group and click **Edit**.





Step 4 (Active/Active mode only) To change the preferred role of the failover group, click either **Primary** or **Secondary**. If you used the wizard, failover group 1 is assigned to the primary unit, and failover group 2 is assigned to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired

Step 5 (Active/Active mode only) To configure failover group preemption, check the **Preempt after booting with optional delay of** check box.

If one unit boots before the other, then both failover groups become active on that unit, despite the Primary or Secondary setting. This option causes the failover group to become active on the designated unit automatically when that unit becomes available.

You can enter an optional delay value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit. Valid values are from 1 to 1200.

Note If Stateful Failover is enabled, the preemption is delayed until the connections are replicated from the unit on which the failover group is currently active.

Step 6 To configure the Interface Policy, choose one of the following:

- Number of failed interfaces that triggers failover—Define a specific number of interfaces that must fail to trigger failover, from 1 to 250. When the number of failed monitored interfaces exceeds the value you specify, the ASA fails over.
- Percentage of failed interfaces that triggers failover—Define a percentage of configured interfaces that must fail to trigger failover. When the number of failed monitored interfaces exceeds the percentage you set, the ASA fails over.

Note Do not use the "Use system failover interface policy" option. You can only set the policy per group at this time.

Step 7 For Active/Standby mode, configure interface poll times in the Failover Poll Time area.

For Active/Active mode, configure interface poll times on the Add/Edit Failover Group dialog box.

- 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.

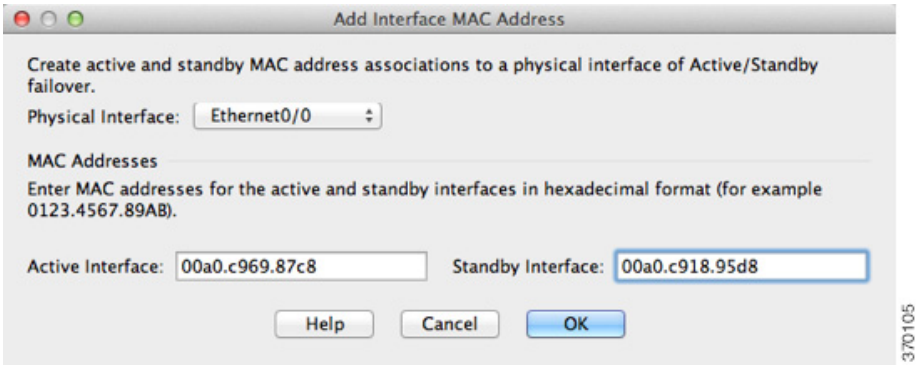
Step 8 (Active/Active mode only) To enable HTTP replication, check the **Enable HTTP replication** check box. For Active/Standby mode, see the "Modifying the Failover Setup" section on page 9-47. For both modes, see the "Modifying the Failover Setup" section for the HTTP replication rate.

Step 9 For Active/Standby mode, to configure virtual MAC addresses, click the **MAC Addresses** tab.

For Active/Active mode, go to the bottom of the Active/Active tab.

You can also set the MAC address using other methods, but we recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.

Step 10 To add a new virtual MAC address entry, click **Add**.



The Add/Edit Interface MAC Address dialog box appears.

Step 11 Choose and interface from the Physical Interface drop-down list.

Step 12 In the Active MAC Address field, type the new MAC address for the active interface.

Step 13 In the Standby MAC Address field, type the new MAC address for the standby interface.

Step 14 Click **OK**.

The interface is added to the table.

Step 15 (Active/Active mode only) Click **OK**.

Step 16 Click **Apply**.

Step 17 (Active/Active mode only) Repeat this procedure for the other failover group, if desired.

Configuring Interface Monitoring and Standby Addresses

By default, monitoring is enabled on all physical interfaces, or for the ASA 5505 and ASASM, all VLAN interfaces. You might want to exclude interfaces attached to less critical networks from affecting your failover policy.

If you did not configure the standby IP addresses in the wizard, you can configure them manually.

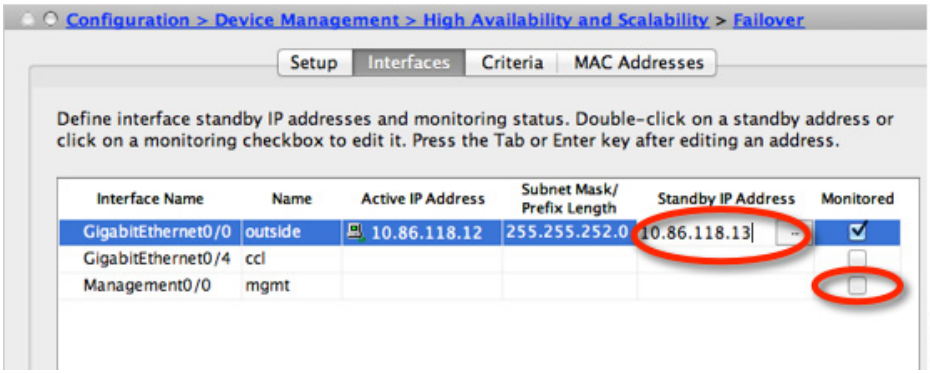
Guidelines

- You can monitor up to 250 interfaces on a unit (across all contexts in multiple context mode).
- In multiple context mode, configure interfaces within each context.

Detailed Steps

Step 1 In single mode, choose **Configuration > Device Management > High Availability > Failover > Interfaces**.

In multiple context mode, within a context choose **Configuration > Device Management > Failover > Interfaces**



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.

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 link and state link do not display IP address; you cannot change those addresses from this tab.

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.

Step 4 For each interface that does not have a standby IP address, double-click the Standby IP Address field and enter an IP address into the field.

Step 5 Click **Apply**.

Configuring Support for Asymmetrically Routed Packets (Active/Active Mode)

When running in Active/Active failover, a unit may receive a return packet for a connection that originated through its peer unit. Because the ASA that receives the packet does not have any connection information for the packet, the packet is dropped. This drop most commonly occurs when the two ASAs in an Active/Active failover pair are connected to different service providers and the outbound connection does not use a NAT address.

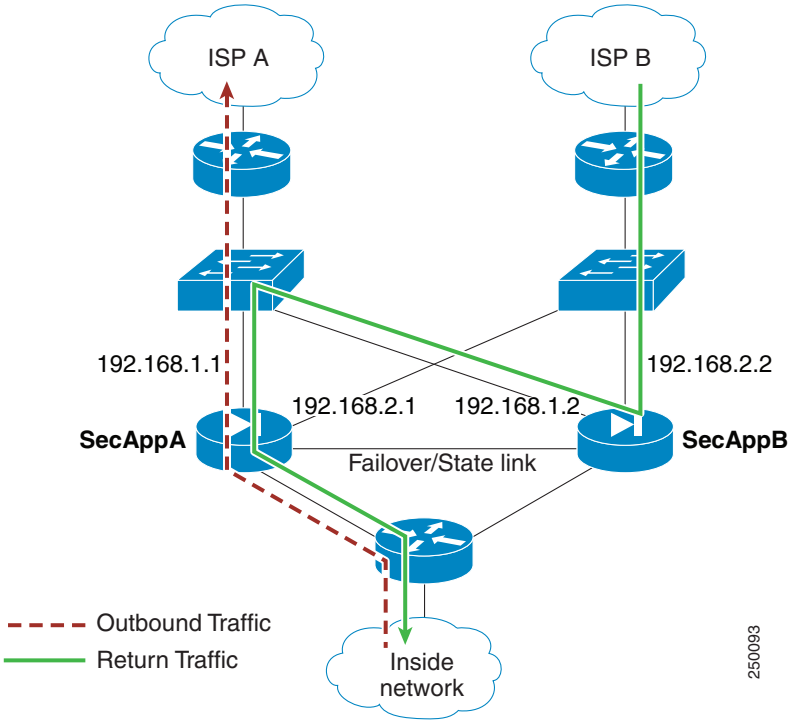
You can prevent the return packets from being dropped by allowing asymmetrically routed packets. To do so, you assign the similar interfaces on each ASA to the same ASR group. For example, both ASAs connect to the inside network on the inside interface, but connect to separate ISPs on the outside interface. On the primary unit, assign the active context outside interface to ASR group 1; on the secondary unit, assign the active context outside interface to the same ASR group 1. When the primary unit outside interface receives a packet for which it has no session information, it checks the session information for the other interfaces in standby contexts that are in the same group; in this case, ASR group 1. If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all of the layer 2 header is rewritten and the packet is reinjected into the stream.

Note This feature does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

Figure 9-13 shows an example of an asymmetrically routed packet.

Figure 9-13 ASR Example



1. An outbound session passes through the ASA with the active SecAppA context. It exits interface outsideISP-A (192.168.1.1).
2. Because of asymmetric routing configured somewhere upstream, the return traffic comes back through the interface outsideISP-B (192.168.2.2) on the ASAwith the active SecAppB context.
3. Normally the return traffic would be dropped because there is no session information for the traffic on interface 192.168.2.2. However, the interface is configured as part of ASR group 1. The unit looks for the session on any other interface configured with the same ASR group ID.
4. The session information is found on interface outsideISP-A (192.168.1.2), which is in the standby state on the unit with SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is rewritten with information for interface 192.168.1.1 and the traffic is redirected out of the interface 192.168.1.2, where it can then return through the interface on the unit from which it originated (192.168.1.1 on SecAppA). This forwarding continues as needed until the session ends.

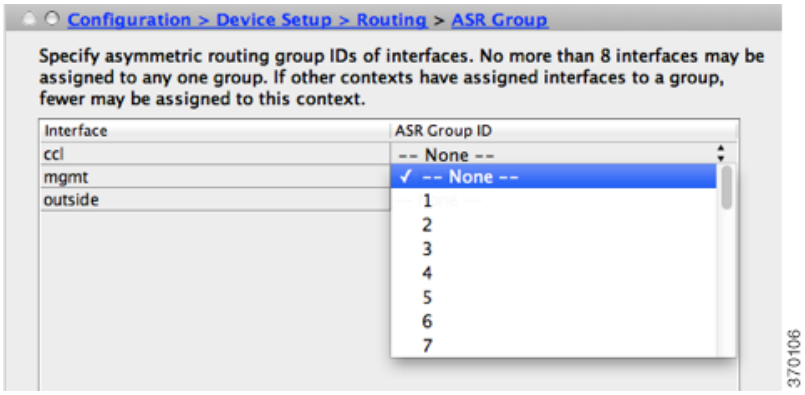
Prerequisites

- Stateful Failover—Passes state information for sessions on interfaces in the active failover group to the standby failover group.

- Replication HTTP—HTTP session state information is not passed to the standby failover group, and therefore is not present on the standby interface. For the ASA to be able to re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.
- Perform this procedure within each active context on the primary and secondary units.

Detailed Steps

Step 1 On the primary unit active context, choose **Configuration > Device Setup > Routing > ASR Groups**.



Step 2 For the interface that receives asymmetrically routed packets, choose an ASR group number from the drop-down list.

Step 3 Click **Apply** to save your changes to the running configuration.

Step 4 Connect ASDM to the secondary unit, and choose the active context similar to the primary unit context.

Step 5 Choose **Configuration > Device Setup > Routing > ASR Groups**.

Step 6 For the similar interface on this unit, choose the same ASR group number.

Step 7 Click **Apply** to save your changes to the running configuration.

Managing Failover

- Modifying the Failover Setup, page 9-47
- Forcing Failover, page 9-50
- Disabling Failover, page 9-51
- Restoring a Failed Unit, page 9-51
- Re-Syncing the Configuration, page 9-52

Modifying the Failover Setup

If you do not use the wizard, or want to change a setting, you can configure the failover setup manually. This section also includes the following options that are not included in the wizard, so you must configure them manually:

- IPsec preshared key for encrypting failover traffic
- HTTP replication rate
- HTTP replication (Active/Standby mode)

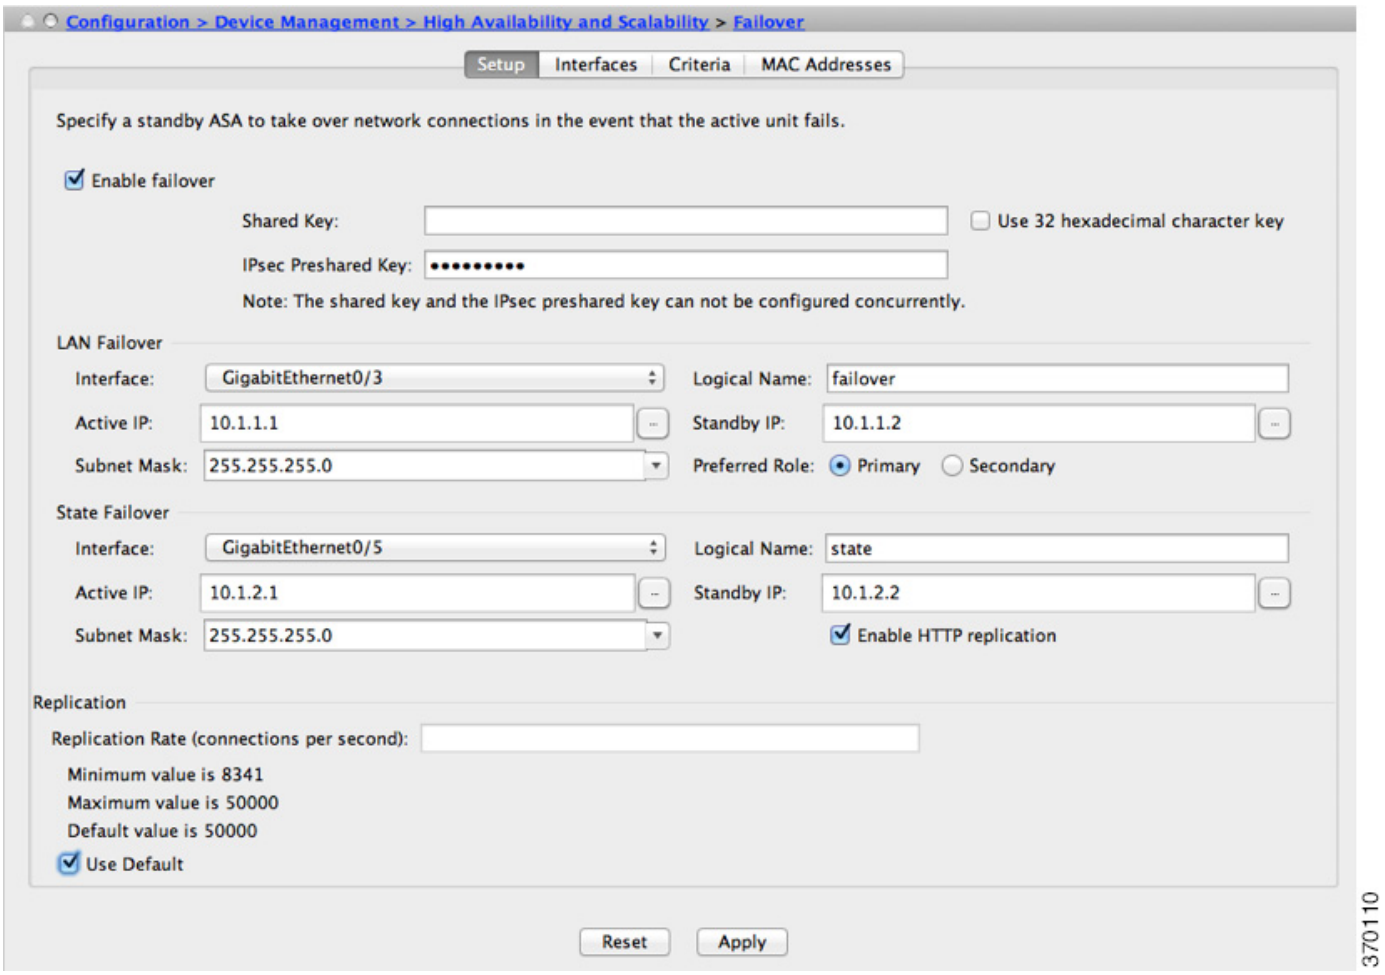
Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Step 1 In single mode, choose **Configuration > Device Management > High Availability and Scalability > Failover > Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.



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 communications on the failover and state links, use one of the following options:

- IPsec Preshared Key (preferred)—The preshared key is used by IKEv2 to establish IPsec LAN-to-LAN tunnels on the failover links between the failover units. Note: failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.
- Secret Key—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.

  Use 32 hexadecimal character key—To use a 32-hexadecimal key for the secret key, check this check box.

Step 4 In the LAN Failover area, set the following parameters for the failover link:

- Interface—Choose the interface to use for the failover link. 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 failover link. Once you specify an interface as the failover link, you cannot edit that interface in the Configuration > Interfaces pane.
- Logical Name—Specify the logical name of the interface used for failover communication, such as “failover”. This name is informational.
- Active IP—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet.
- Standby IP—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- Subnet Mask—Specify the subnet mask.
- Preferred Role—Select **Primary** or **Secondary** to specify whether the preferred role for this ASA is as the primary or secondary unit.

Step 5 (Optional) Configure the state link by doing the following:

- Interface—Choose the interface to use for the state link. You can choose an unconfigured interface or subinterface, the failover link, or the **--Use Named--** option.

  Note We recommend that you use two separate, dedicated interfaces for the failover link and the state link.

  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 failover link, you do not need to specify the Active IP, Subnet Mask, Logical Name, and Standby IP values; the values specified for the failover link 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.
- Logical Name—Specify the logical name of the interface used for state communication, such as “state”. This name is informational.
- Active IP—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet, different from the failover link.
- Standby IP—Specify the standby IP address for the interface, on the same subnet as the active IP address.

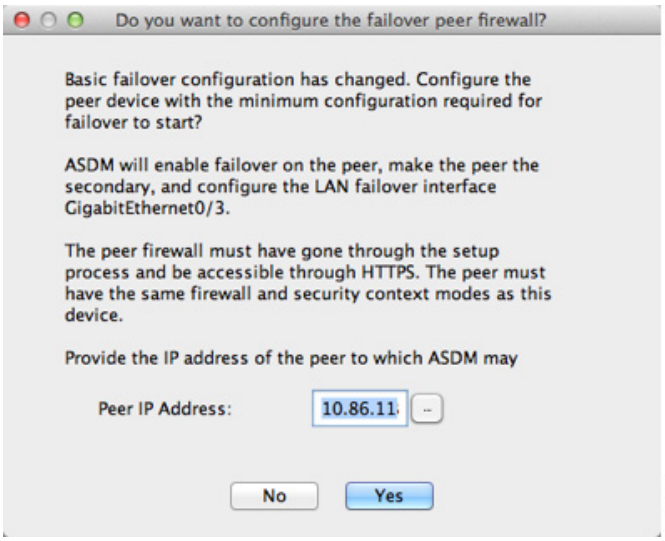
- Subnet Mask—Specify the subnet mask.
- (Optional, Active/Standby only) Enable HTTP Replication—Enable HTTP replication by checking the **Enable HTTP Replication** check box. This option enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover. In Active/Active mode, set the HTTP replication per failover group. See the "Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses" section on page 9-41.

Step 6 In the Replication area, set the HTTP replication rate between 8341 connections per second and 50000. The default is 50000. To use the default, check the **Use Default check** box.

Step 7 Click **Apply**.

The configuration is saved to the device.

Step 8 If you are enabling failover, you see a dialog box to configure the failover peer.



- Click **No** if you want to connect to the failover peer later and configure the matching settings manually.
- Click **Yes** to let ASDM automatically configure the relevant failover settings on the failover peer. Provide the peer IP address in the Peer IP Address field.

Forcing Failover

To force the standby unit to become active, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Step 1 To force failover at the unit level:

a. Choose the screen depending on your context mode:

- In single context mode choose **Monitoring > Properties > Failover > Status**.
- In multiple context mode, in the System choose **Monitoring > Failover > System**.

b. Click one of the following buttons:

- Click **Make Active** to make the unit this unit.
- Click **Make Standby** to make the other unit the active unit.

Step 2 (Active/Active mode only) To force failover at the failover group level:

a. In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

b. Click one of the following buttons:

- Click **Make Active** to make the failover group active on this unit.
- Click **Make Standby** to make the failover group active on the other unit.

Disabling Failover

To disable failover, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Step 1 In single mode, choose **Configuration > Device Management > High Availability and Scalability > Failover > Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.

Step 2 Uncheck the **Enable Failover** check box.

Step 3 Click **Apply**.

Restoring a Failed Unit

To restore a failed unit to an unfailed state, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Step 1 To restore failover at the unit level:

a. Choose the screen depending on your context mode:

- In single context mode choose **Monitoring > Properties > Failover > Status**.
- In multiple context mode, in the System choose **Monitoring > Failover > System**.

b. Click **Reset Failover**.

Step 2 (Active/Active mode only) To reset failover at the failover group level:

a. In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

b. Click **Reset Failover**.

Re-Syncing 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**.

Monitoring Failover

- Failover Messages, page 9-52
- Monitoring Failover, page 9-53

Failover Messages

When a failover occurs, both ASAs send out system messages. This section includes the following topics:

- Failover Syslog Messages, page 9-52
- Failover Debug Messages, page 9-53
- SNMP Failover Traps, page 9-53

Failover Syslog Messages

The ASA issues a number of syslog messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the syslog messages guide. To enable logging, see Chapter 41, “Configuring Logging.”

Note During a fail over, failover logically shuts down and then bring up interfaces, generating syslog messages 411001 and 411002. This is normal activity.

Failover 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 Failover Traps

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 42, "Configuring SNMP" for more information.

Monitoring 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.

Use the following screens in the Monitoring > Properties > Failover area to monitor Active/Active failover:

- System, page 9-53
- Failover Group 1 and Failover Group 2, page 9-54

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 Failover

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

Table 9-4 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
Active/Standby failover	7.0(1)	This feature was introduced.
Active/Active failover	7.0(1)	This feature was introduced.
Support for a hex value for the failover key	7.0(4)	You can now specify a hex value for failover link encryption. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.

Table 9-4 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
Support for the master passphrase for the failover key	8.3(1)	The failover key now supports the master passphrase, which encrypts the shared key in the running and startup configuration. If you are copying the shared secret from one ASA to another, for example from the **more system:running-config** command, you can successfully copy and paste the encrypted shared key. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. There were no ASDM changes.
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
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	9.1(2)	Instead of using the proprietary encryption for the failover key, you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.



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.

Note Some features are not supported when using clustering. See the "Unsupported Features" section on page 10-20.

- Information About ASA Clustering, page 10-1
- Licensing Requirements for ASA Clustering, page 10-27
- Prerequisites for ASA Clustering, page 10-27
- Guidelines and Limitations, page 10-28
- Default Settings, page 10-32
- Configuring ASA Clustering, page 10-32
- Managing ASA Cluster Members, page 10-59
- Monitoring the ASA Cluster, page 10-70
- Configuration Examples for ASA Clustering, page 10-72
- Feature History for ASA Clustering, page 10-88

Information About ASA Clustering

- 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 10-9
- Configuration Replication, page 10-11
- ASA Cluster Management, page 10-11
- Load Balancing Methods, page 10-13

- Inter-Site Clustering, page 10-16
- How the ASA Cluster Manages Connections, page 10-18
- ASA Features and Clustering, page 10-20

How the ASA Cluster Fits into Your Network

The cluster consists of multiple ASAs acting as a single unit. (See the "Licensing Requirements for ASA Clustering" section on page 10-27 for the number of units supported per model). 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 10-11.

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 10-13.
- Policy-Based Routing (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using route maps and ACLs. See the "Policy-Based Routing (Routed Firewall Mode Only)" section on page 10-15.
- Equal-Cost Multi-Path Routing (Routed firewall 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 Firewall Mode Only)" section on page 10-16.

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

- ASA Hardware and Software Requirements, page 10-3
- Bootstrap Configuration, page 10-3
- Master and Slave Unit Roles, page 10-3
- Master Unit Election, page 10-3

ASA Hardware and Software Requirements

All units in a cluster:

- Must be the same model with the same DRAM. You do not have to have the same amount of flash memory.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported. See the "Upgrade Path and Migrations" section on page 46-1.
- (9.1.3 and Earlier) Must be in the same geographical location. (9.1(4) and later) You can have cluster members in different geographical locations (inter-site) when using individual interface mode. See the "Inter-Site Clustering" section on page 10-16 for more information.
- Must be in the same security context mode, single or multiple.
- (Single context mode) Must be in the same firewall mode, routed or transparent.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the master unit for initial cluster control link communication before configuration replication.
- Must have the same cluster, encryption and, for the ASA 5585-X, 10 GE I/O licenses.

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 10-21.

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 10-21 for a list of centralized features.

ASA Cluster Interfaces

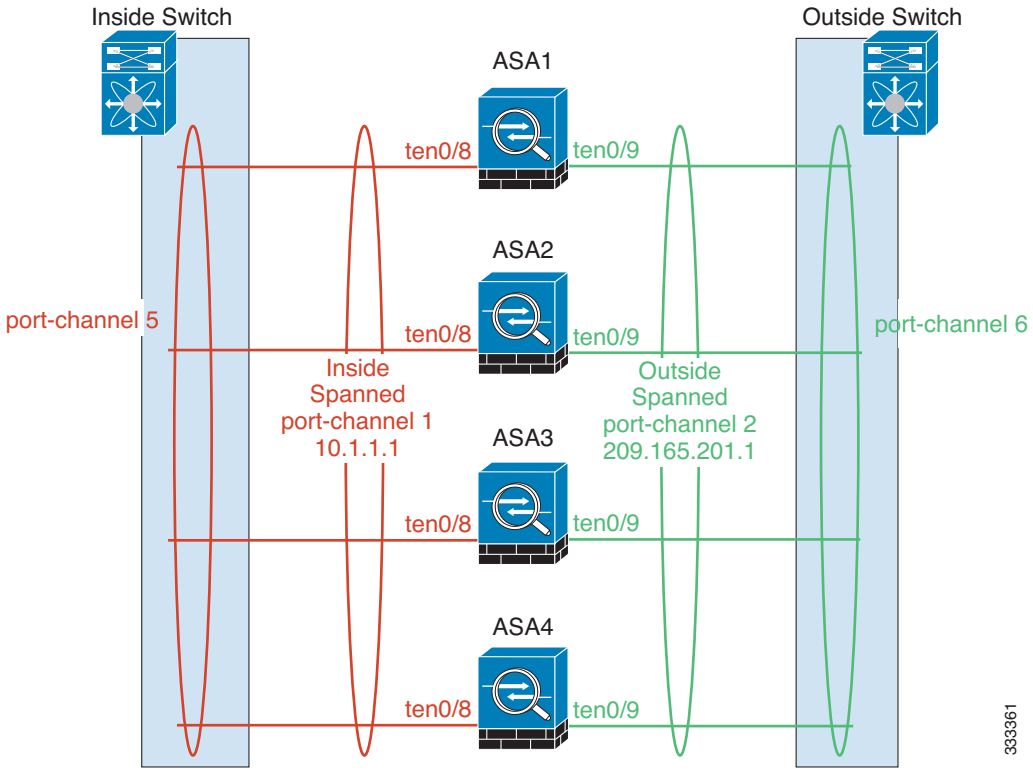
You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. All data interfaces in the cluster must be one type only.

- Interface Types, page 10-4
- Interface Type Mode, page 10-6

Interface Types

- Spanned EtherChannel (Recommended)

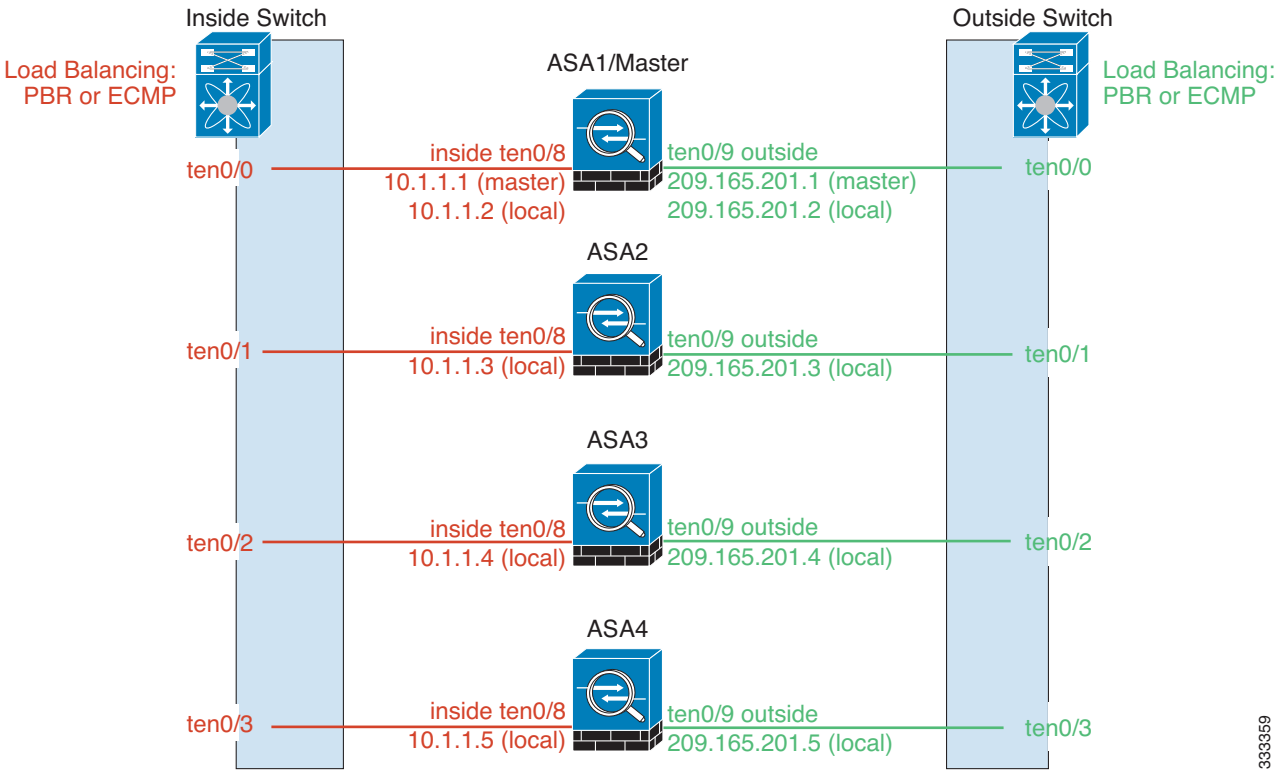
  You can group one or more interfaces per unit into an EtherChannel that spans all units in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed 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 10-13.



- Individual interfaces (Routed firewall 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 one for 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 10-13.

Note We recommend Spanned EtherChannels instead of Individual interfaces because Individual interfaces rely on routing protocols to load-balance traffic, and routing protocols often have slow convergence during a link failure.



Interface Type Mode

You must choose the interface type (Spanned EtherChannel or Individual) 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 10-7
- Cluster Control Link Interfaces and Network, page 10-7
- Sizing the Cluster Control Link, page 10-7
- Cluster Control Link Redundancy, page 10-8

- Cluster Control Link Reliability, page 10-8
- Cluster Control Link Failure, page 10-9

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 10-11.)
- Health monitoring. (See the "Unit Health Monitoring" section on page 10-9.)

Data traffic includes:

- State replication. (See the "Data Path Connection State Replication" section on page 10-10.)
- Connection ownership queries and data packet forwarding. (See the "Rebalancing New TCP Connections Across the Cluster" section on page 10-20.)

Cluster Control Link Interfaces and Network

You can use any data interface(s) for the cluster control link, with the following exceptions:

- 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 ASA IPS module, you cannot use the module interfaces for the cluster control link; you can, however, use interfaces on the ASA 5585-X Network Module.

You can use an EtherChannel or redundant interface; see the "Cluster Control Link Redundancy" section on page 10-8 for more information.

For the ASA 5585-X with SSP-10 and SSP-20, which include 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 10-7 for more information.

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.

For a 2-member cluster, do not directly-connect the cluster control link from one ASA to the other ASA. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Sizing the Cluster Control Link

You should size the cluster control link to match the expected throughput of each member. 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 at least 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 feature, 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.

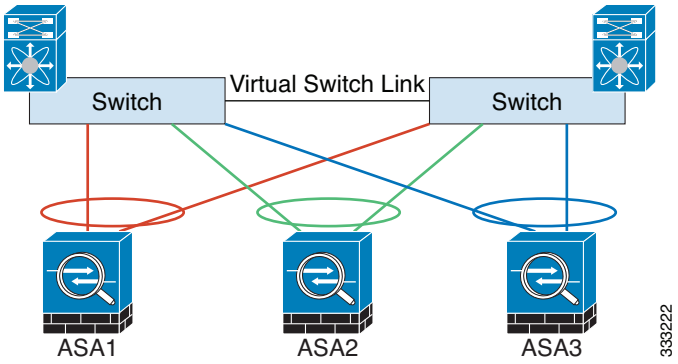
Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

(9.1(4) and later) For inter-site clusters and sizing the data center interconnect for cluster control link traffic, see the "Inter-Site Clustering" section on page 10-16.

Cluster Control Link Redundancy

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 Reliability

(9.1(4) and later) To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

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. After you fix the cluster control link, you must manually rejoin the cluster by re-enabling clustering; see Rejoining the Cluster, page 10-10.

Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster, the management interface is not accessible (because it then uses the Main IP address, which is the same as the master unit). You must use the console port for any further configuration.

High Availability within the ASA Cluster

- Unit Health Monitoring, page 10-9
- Interface monitoring, page 10-9
- Unit or Interface Failure, page 10-9
- Data Path Connection State Replication, page 10-10

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. The amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. For EtherChannels (spanned or not), if the interface is down on an established member, then the ASA

removes the member after 9 seconds. If the unit is joining the cluster as a new member, the ASA waits 45 seconds before rejecting the new unit. For non-EtherChannels, the unit is removed after 500 ms, regardless of the member state.

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.

The ASA automatically tries to rejoin the cluster; see Rejoining the Cluster, page 10-10.

Note When an ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster, the management interface is not accessible (because it then uses the Main IP address, which is the same as the master unit). You must use the console port for any further configuration.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering according to Configuring ASA Cluster Parameters, page 10-59.
- Failed data interface—The ASA automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the ASA disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering according to Configuring ASA Cluster Parameters, page 10-59.
- Failed unit—If the unit was removed from the cluster because of a unit health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the unit will rejoin the cluster when it starts up again as long as the cluster control link is up and clustering is still enabled.

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 kind of 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	
SNMP Engine ID	**No**	
VPN (Site-to-Site)	**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 10-11
- Management Interface, page 10-11
- Master Unit Management Vs. Slave Unit Management, page 10-12
- RSA Key Replication, page 10-12
- ASDM Connection Certificate IP Address Mismatch, page 10-12

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, including the master unit, uses the Local IP address to connect to the server.

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. Note that 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.
- SNMP
- NetFlow

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 40, "Configuring Digital Certificates."

Load Balancing Methods

See also the "ASA Cluster Interfaces" section on page 10-4.

- Spanned EtherChannel (Recommended), page 10-13
- Policy-Based Routing (Routed Firewall Mode Only), page 10-15
- Equal-Cost Multi-Path Routing (Routed Firewall Mode Only), page 10-16

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 10-13
- Guidelines for Maximum Throughput, page 10-13
- Load Balancing, page 10-14
- EtherChannel Redundancy, page 10-14
- Connecting to a VSS or vPC, page 10-14

Spanned EtherChannel Benefits

The EtherChannel method of load-balancing is recommended over other methods for the following benefits:

- Faster failure discovery.
- Faster convergence time. Individual interfaces rely on routing protocols to load-balance traffic, and routing protocols often have slow convergence during a link failure.
- Ease of configuration.

For more information about EtherChannels in general (not just for clustering), see the "EtherChannels" section on page 11-5.

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 IP addresses and TCP and UDP port numbers.

Note On the ASA, do not change the load-balancing algorithm from the default (see the "Customizing the EtherChannel" section on page 11-33). On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Nexus OS or IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the ASAs in a cluster.

The number of links in the EtherChannel affects load balancing. See the "Load Balancing" section on page 11-7 for more information.

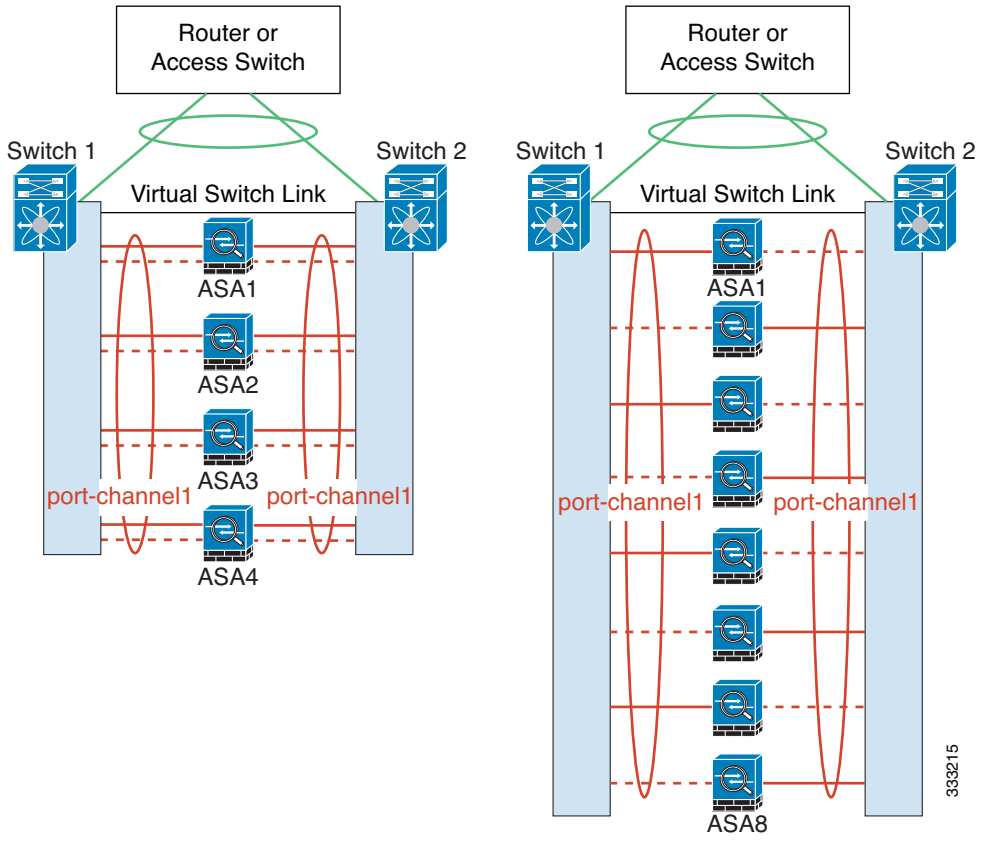
Symmetric load balancing is not always possible. If you configure NAT, then forward and return packets will have different IP addresses and/or ports. Return traffic will be sent to a different unit based on the hash, and the cluster will have to redirect most returning traffic to the correct unit. See the "NAT" section on page 10-24 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 Firewall Mode Only)

When using Individual interfaces, each ASA interface maintains its own IP address and MAC address. One method of load balancing 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 Firewall Mode Only)

When using Individual interfaces, each ASA interface maintains its own IP address and MAC address. One method of load balancing 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.

Inter-Site Clustering

9.1(4) and later

- Inter-Site Clustering Guidelines, page 10-16
- Sizing the Data Center Interconnect, page 10-17
- Inter-Site Example, page 10-17

Inter-Site Clustering Guidelines

See the following guidelines for inter-site clustering:

- Individual Interface mode only.
- The cluster control link latency must be less than 20 ms round-trip time (RTT).
- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing (see the "Rebalancing New TCP Connections Across the Cluster" section on page 10-20); you do not want connections rebalanced to cluster members at a different site.
- The cluster implementation does not differentiate between members at multiple sites; therefore, connection roles for a given connection may span across sites (see the "Connection Roles" section on page 10-18). This is expected behavior.

Sizing the Data Center Interconnect

You should reserve bandwidth on the data center interconnect (DCI) for cluster control link traffic equivalent to the following calculation:

$$\frac{\text{\# of cluster members per site}}{2} \times \text{cluster control link size per member}$$

If the number of members differs at each site, use the larger number for your calculation. The minimum bandwidth for the DCI should not be less than the size of the cluster control link for one member.

For example:

- For 4 members at 2 sites:
  - 4 cluster members total
  - 2 members at each site
  - 5 Gbps cluster control link per member

  Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).
- For 8 members at 2 sites, the size increases:
  - 8 cluster members total
  - 4 members at each site
  - 5 Gbps cluster control link per member

  Reserved DCI bandwidth = 10 Gbps (4/2 x 5 Gbps).
- For 6 members at 3 sites:
  - 6 cluster members total
  - 3 members at site 1, 2 members at site 2, and 1 member at site 3
  - 10 Gbps cluster control link per member

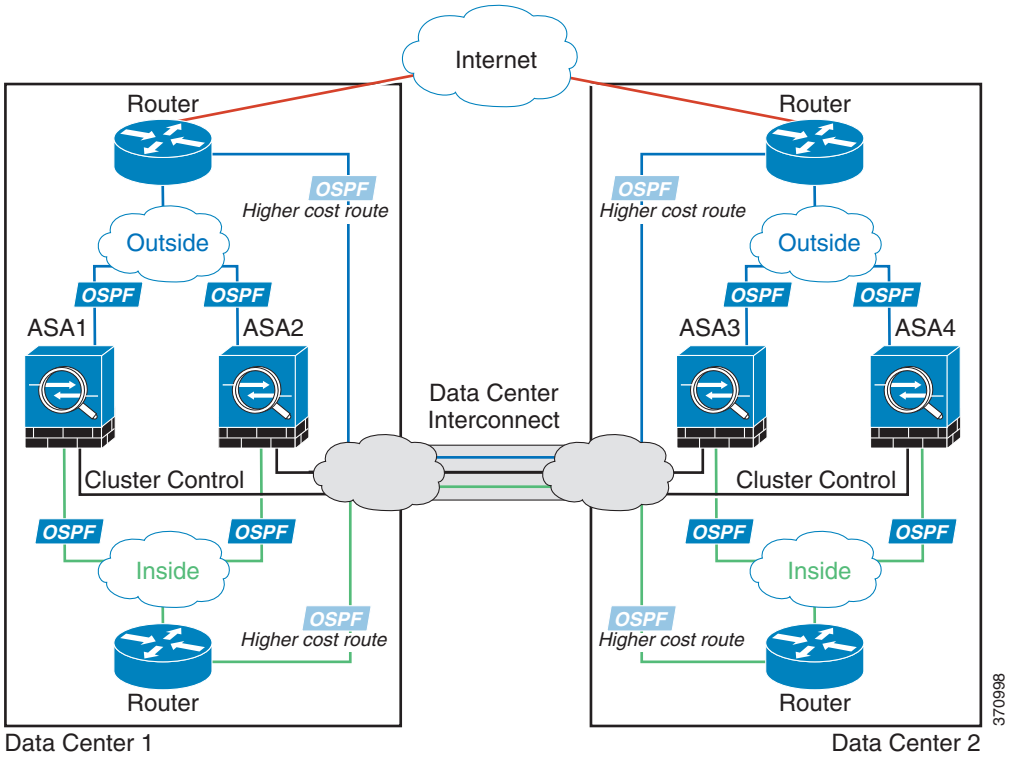
  Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).
- For 2 members at 2 sites:
  - 2 cluster members total
  - 1 member at each site
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 10 Gbps (1/2 x 10 Gbps = 5 Gbps; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).

Inter-Site Example

The following example shows 2 ASA cluster members at each of 2 data centers. The cluster members are connected by the cluster control link over the DCI. The inside and outside routers at each data center use OSPF and PBR or ECMP to load balance the traffic between cluster members. By assigning a higher

cost route across the DCI, traffic stays within each data center unless all ASA cluster members at a given site go down. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the ASA cluster members at the other site.



How the ASA Cluster Manages Connections

- Connection Roles, page 10-18
- New Connection Ownership, page 10-19
- Sample Data Flow, page 10-19
- Rebalancing New TCP Connections Across the Cluster, page 10-20

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.

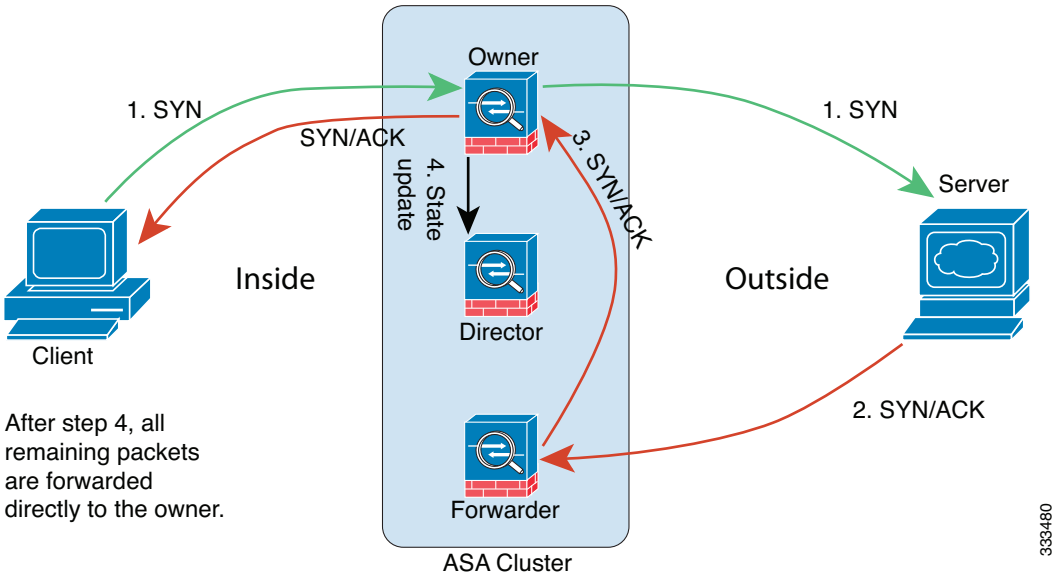
- Forwarder—A unit that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that 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.

New Connection Ownership

When a new connection is directed to a member of the cluster via load balancing, 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 10-13.

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 10-20
- Centralized Features, page 10-21
- Features Applied to Individual Units, page 10-22
- Dynamic Routing, page 10-22
- Multicast Routing, page 10-24
- NAT, page 10-24
- AAA for Network Access, page 10-25
- Syslog and Netflow, page 10-26
- SNMP, page 10-26
- VPN, page 10-26
- FTP, page 10-26
- Cisco TrustSec, page 10-27

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 site-to-site 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 10-7 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 10-20), 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 or to the master unit.

- 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.

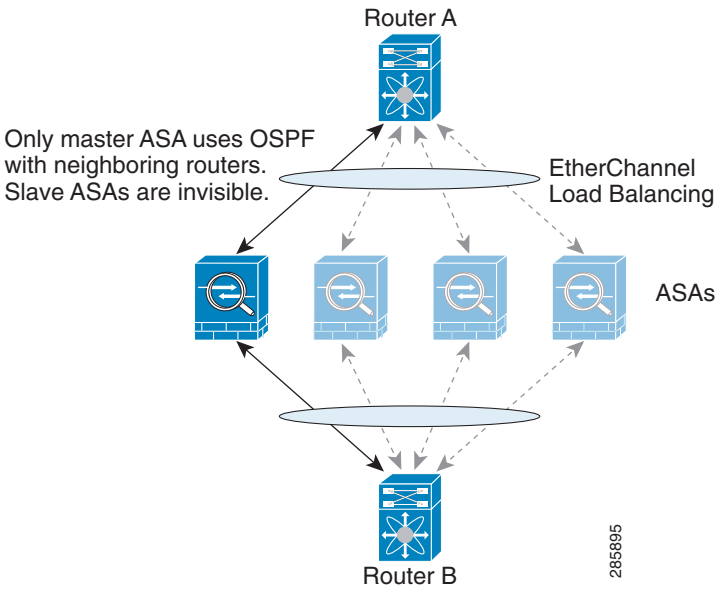
Dynamic Routing

- Dynamic Routing in Spanned EtherChannel Mode, page 10-22
- Dynamic Routing in Individual Interface Mode, page 10-23

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



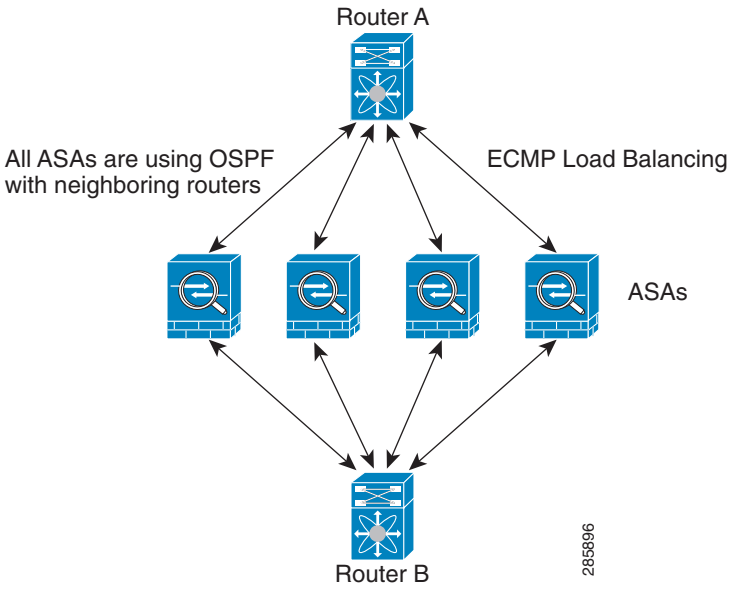
After the slave members learn the routes from the master unit, each unit makes forwarding decisions independently.

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 10-24
- Multicast Routing in Individual Interface Mode, page 10-24

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 for dynamic PAT—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 section in the firewall configuration guide.
- No static PAT for the following inspections—
  - FTP
  - PPTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - All Voice-over-IP applications

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 41-22.
- 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.

Note 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 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.
- If you use AAA for FTP access, then the control channel flow is centralized on the master unit.

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, supports up to 8 units. 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 and the same 10 GE I/O license.
ASA 5512-X[1]	Security Plus license, supports 2 units. **Note** Each unit must have the same encryption license.
ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X[1]	Base License, supports 2 units. **Note** Each unit must have the same encryption license.
All other models	No support.

1. Supported in 9.1(4) and later.

Prerequisites for ASA Clustering

Switch Prerequisites

- Be sure to complete the switch configuration before you configure clustering on the ASAs.
- 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

External Hardware	External Software	ASA Version
Nexus 7000	NXOS 5.2(5) and later	9.0(1) and later.
Nexus 5000	NXOS 7.0(1) and later	9.1(4) and later.
Catalyst 6500 with Supervisor 32, 720, and 720-10GE	IOS 12.2(33)SXI7, SXI8, and SXI9 and later	9.0(1) and later.
Catalyst 3750-X	IOS 15.0(2) and later	9.1(4) and later.

ASA Prerequisites

- Provide each unit with a unique IP address before you join them to the management network.

  - See Chapter 3, "Getting Started," for more information about connecting to the ASA and setting the management IP address.
  - 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.
- To use jumbo frames on the cluster control link (recommended), you must enable Jumbo Frame Reservation before you enable clustering. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39.
- See also the "ASA Hardware and Software Requirements" section on page 10-3.

Other Prerequisites

We recommend using a terminal server to access all cluster member unit console ports. For initial setup, and ongoing management (for example, when a unit goes down), a terminal server is useful for remote management.

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. For single mode, the firewall mode must match on all units.

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 on:

- ASA 5585-X

  For the ASA 5585-X with SSP-10 and SSP-20, which include 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 10-7 for more information.
- ASA 5580
- ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X

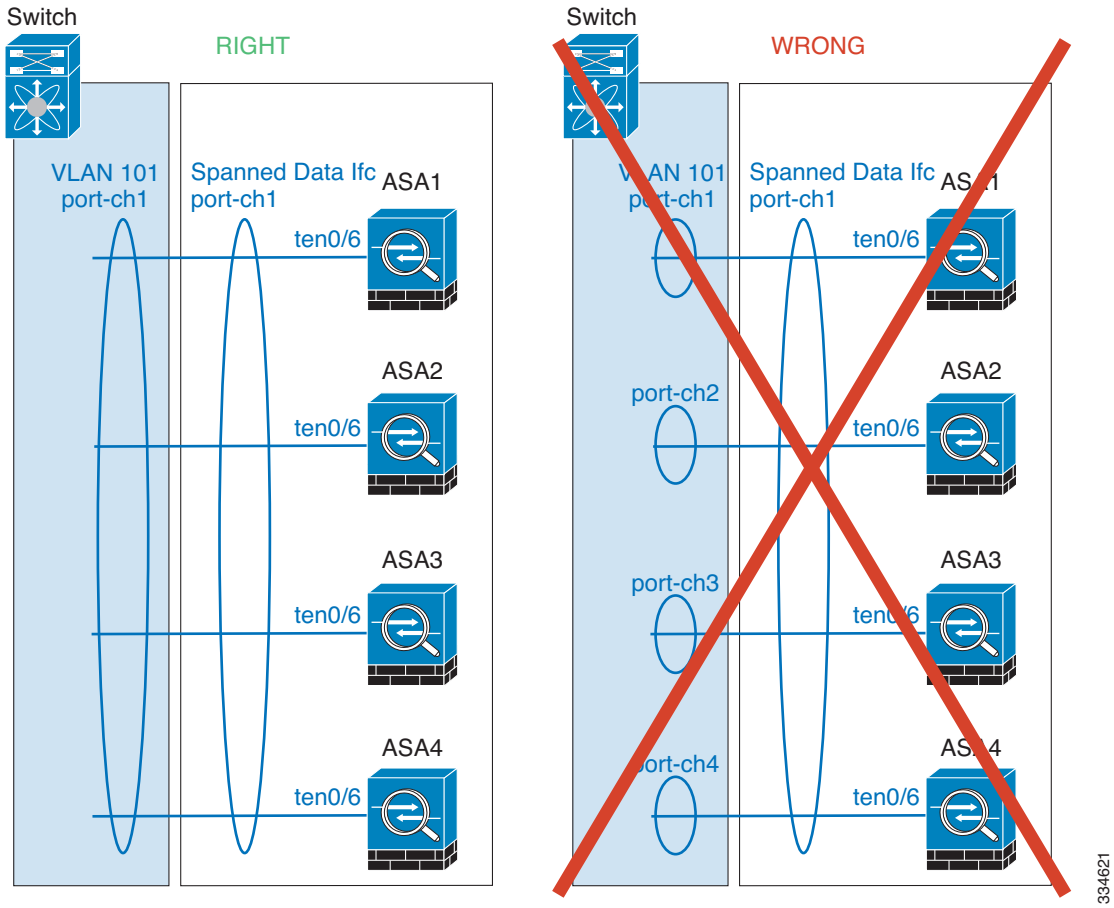
Switch Guidelines

- On the switch(es) 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.

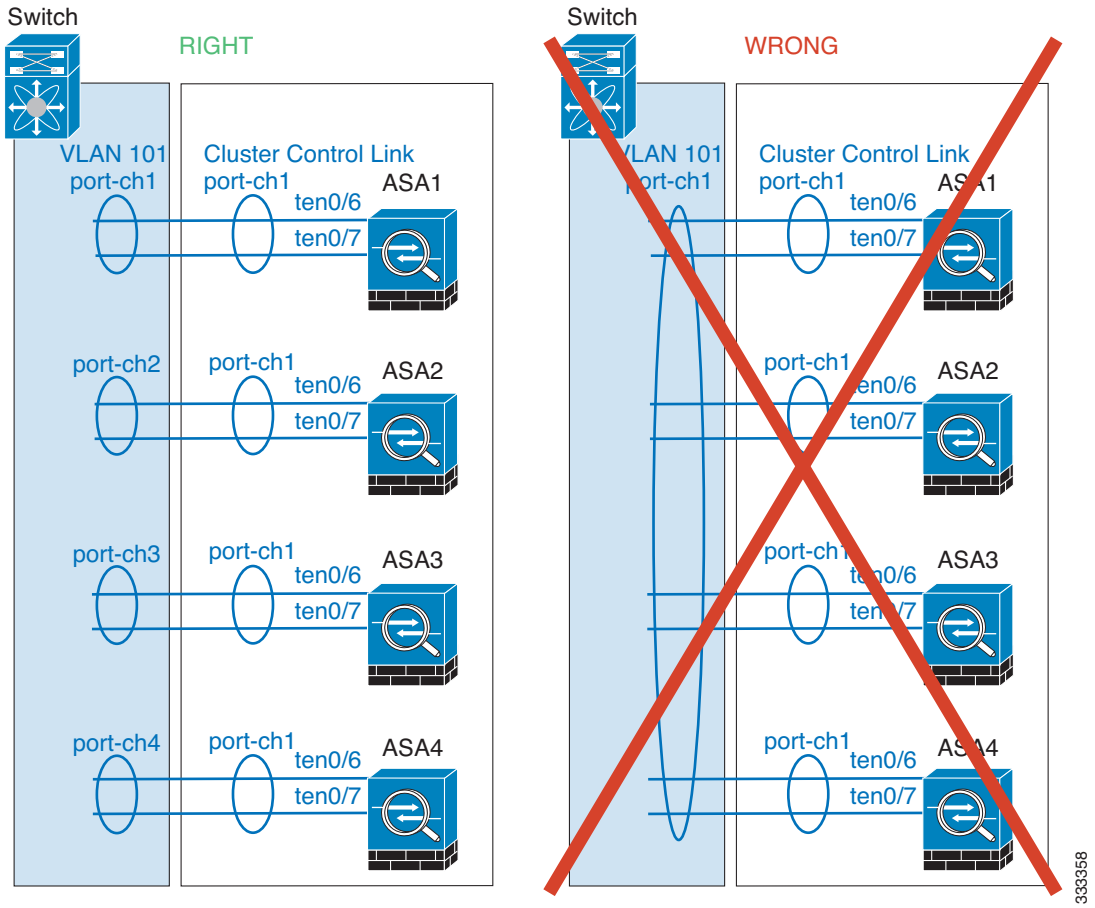
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an Individual interface on the switch.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Nexus OS and IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the ASAs in a cluster. *Do not* change the load-balancing algorithm from the default on the ASA .
- 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.
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Nexus switches.

EtherChannel Guidelines

- 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.
- For detailed EtherChannel guidelines, limitations, and prerequisites, see the "Configuring an EtherChannel" section on page 11-30.
- See also the "EtherChannel Guidelines" section on page 11-13.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - 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.



Additional Guidelines

- See the "ASA Hardware and Software Requirements" section on page 10-3.
- For unsupported features with clustering, see the "Unsupported Features" section on page 10-20.
- When significant 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 use a Windows 2003 server connected to a Spanned EtherChannel, when the syslog server port is down and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the ASA cluster. These messages can result in some units of the ASA cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.

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

Note To enable or disable clustering, you must use a console connection (for CLI) or an ASDM connection.

- Task Flow for ASA Cluster Configuration, page 10-32
- Cabling the Cluster Units and Configuring Upstream and Downstream Equipment, page 10-33
- Backing Up Your Configurations (Recommended), page 10-35
- Configuring the Cluster Interface Mode on the Master Unit, page 10-35
- (Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit, page 10-38
- Adding or Joining an ASA Cluster, page 10-50

Task Flow for ASA Cluster Configuration

To configure clustering, perform the following steps:

Step 1 Complete all pre-configuration on the switches and ASAs according to the "Prerequisites for ASA Clustering" section on page 10-27.

Step 2 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 10-33.

Step 3 (Recommended) Back up each unit configuration before you enable clustering. See the "Backing Up Your Configurations (Recommended)" section on page 10-35.

Step 4 Configure the interface mode. 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 10-35.

Step 5 (Recommended) Configure interfaces for clustering on the master unit. You cannot enable clustering if the interfaces are not cluster-ready. In single context mode, you can alternatively configure many interface settings within the High Availability and Scalability wizard, but not all interface options are available in the wizard, and you cannot configure the interfaces in multiple context mode within the wizard. See the "(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit" section on page 10-38.

Step 6 Join the cluster by running the High Availability and Scalability wizard. See the "Adding or Joining an ASA Cluster" section on page 10-50.

Step 7 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 10-20.

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.

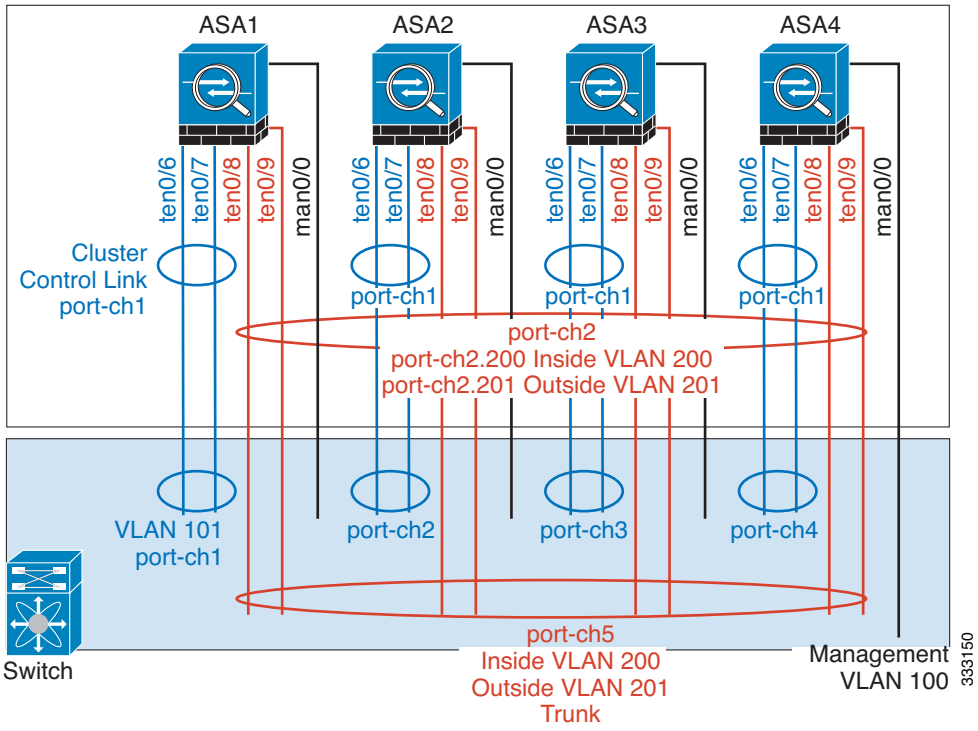
Examples

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). On the switch, configure these VLANs and networks now; 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

Back up your configuration. See the "Backing Up Your Configurations (Recommended)" section on page 10-35.

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 10-66 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 46-21 for a detailed procedure.

What to Do Next

Configure the cluster interface mode on the master unit. See the "Configuring the Cluster Interface Mode on the Master Unit" section on page 10-35.

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 10-6.

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

- 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.

Detailed Steps

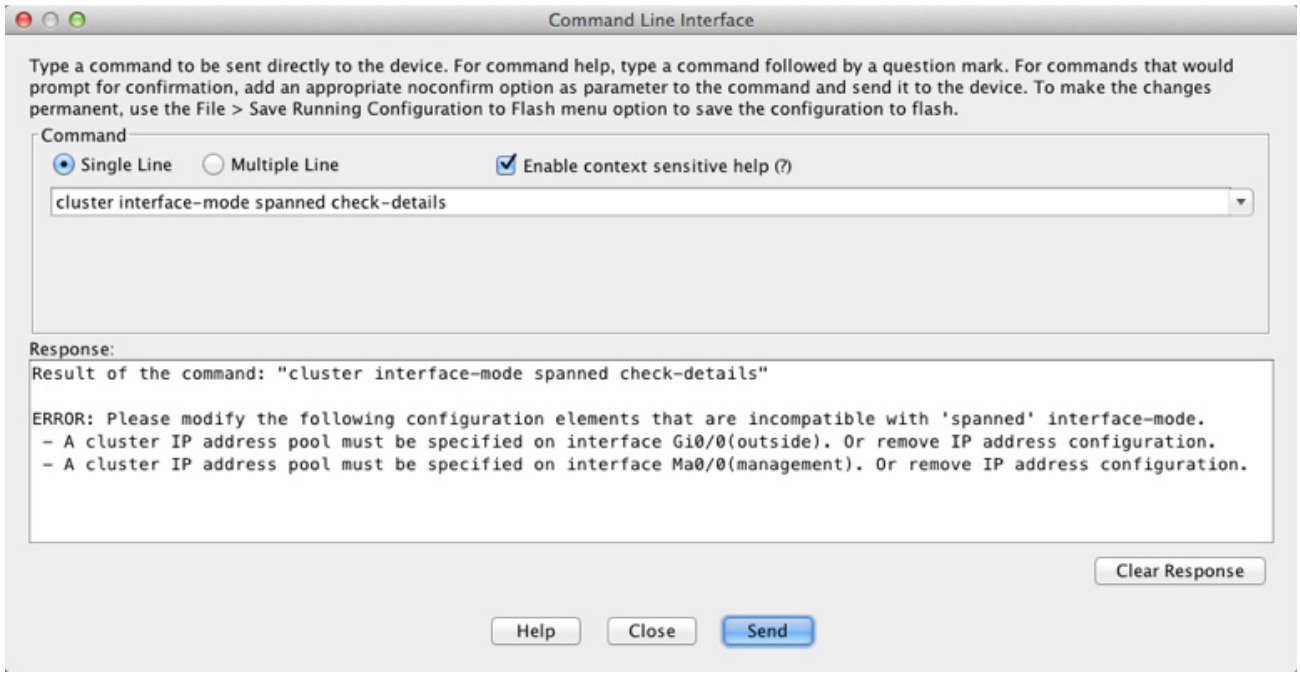
Step 1 In ASDM on the master unit, choose **Tools > Command Line Interface**.

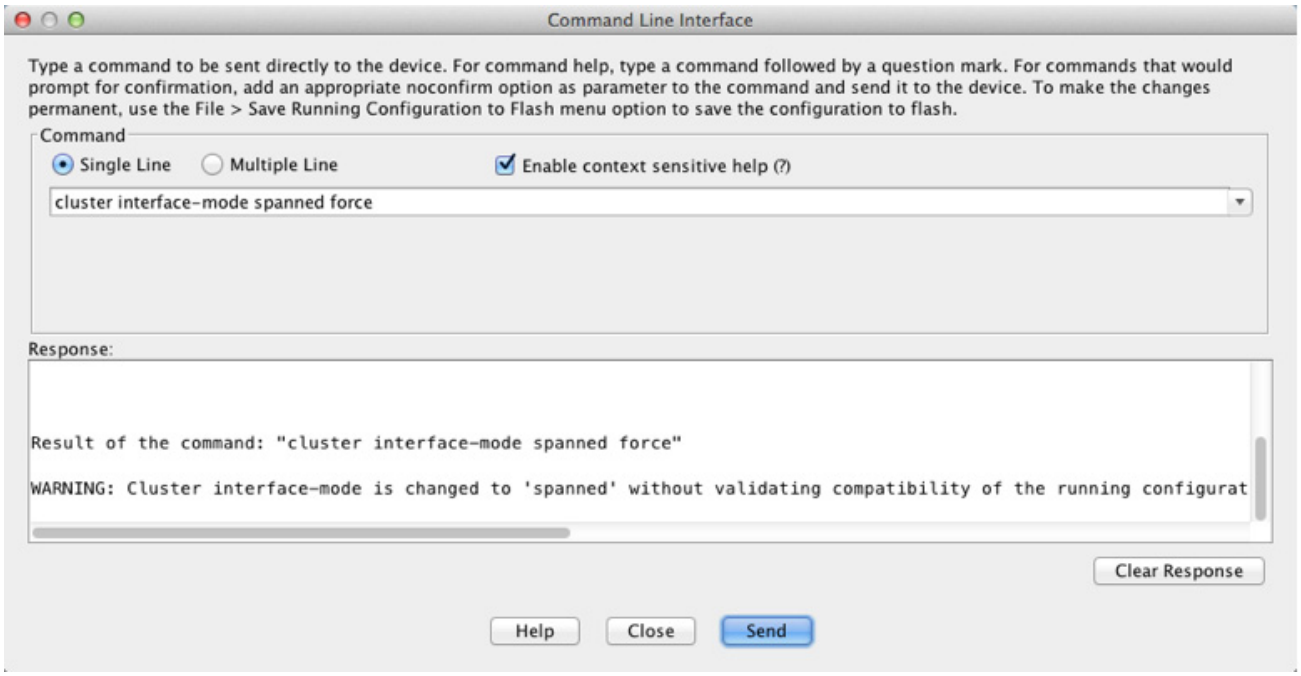
Step 2 Enter the following commands:

Caution After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool), you will not be able to reconnect because cluster-incompatible interface configuration is removed. In that case, you will have to connect to the console port to fix the interface configuration. To configure interfaces to be compatible with clustering, see the "(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit" section on page 10-38.

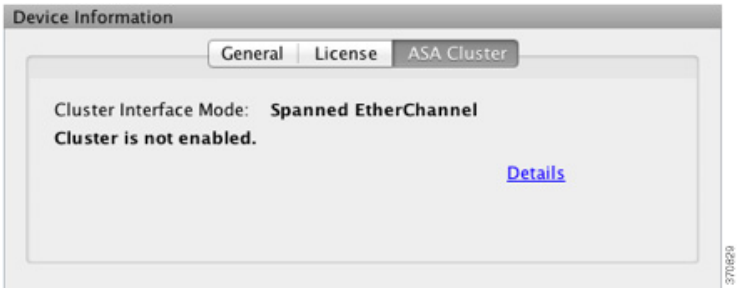
	Command	Purpose
Step 1	**cluster interface-mode** {**individual** | **spanned**} **check-details** **Example:** `cluster interface-mode spanned check-details`	The **check-details** command shows any incompatible configuration so you can force the interface mode and fix your configuration later; the mode is not changed with this command.
Step 2	**cluster interface-mode** {**individual** | **spanned**} **force** **Example:** `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. 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 existing configuration. 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, thus requiring you to connect to the console port to reconfigure your management access. 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.

For example:





Step 3 Quit ASDM and reload. ASDM needs to be restarted to correctly account for the cluster interface mode. After you reload, you see the ASA Cluster tab on the Home Page:



What to Do Next

Configure interfaces. See the "(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit" section on page 10-38.

(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit

You must modify any interface that is currently configured with an IP address to be cluster-ready *before* you enable clustering. At a minimum, you must modify the management interface to which ASDM is currently connected. For other interfaces, you can configure them before or after you enable clustering; we recommend pre-configuring all of your interfaces so that the complete configuration is synced to new cluster members. In multiple context mode, you must use the procedures in this section to fix existing interfaces or to configure new interfaces. However, in single mode, you can skip this section and configure common interface parameters within the High Availability and Scalability wizard (see the "Adding or Joining an ASA Cluster" section on page 10-50). Note that advanced interface settings such as creating EtherChannels for Individual interfaces are not available in the wizard.

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, with the exception of the management interface, which can be an Individual interface even in Spanned EtherChannel mode. For more information, see the "ASA Cluster Interfaces" section on page 10-4.

- Configuring Individual Interfaces (Recommended for the Management Interface), page 10-38
- Configuring Spanned EtherChannels, page 10-43

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 10-11 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 10-35.
- 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 10-13.
- (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 11-30. 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 11-27. Management-only interfaces cannot be redundant interfaces.
  - For subinterfaces, see the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 11-36.
- 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 for a management interface 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 13-1 for more information.

d. Enable Interface—If the interface is not already enabled, check this check box.

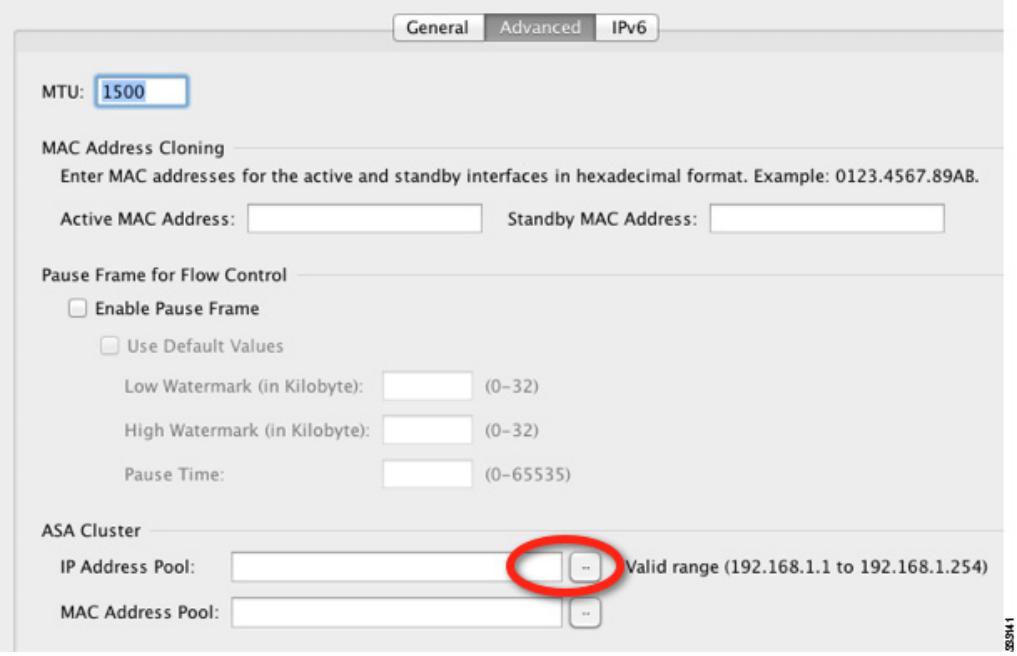
e. Use Static IP—To set the IP address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported.

f. (Optional) Description—Enter a description for this interface. The description can be up to 240 characters on a single line, without carriage returns.

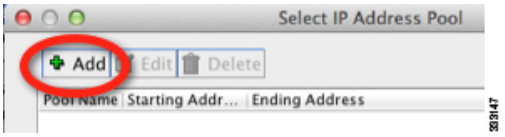
Note For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 11-25.

Step 3 To add the IPv4 cluster IP pool, and optionally a MAC address pool, click the **Advanced** tab.

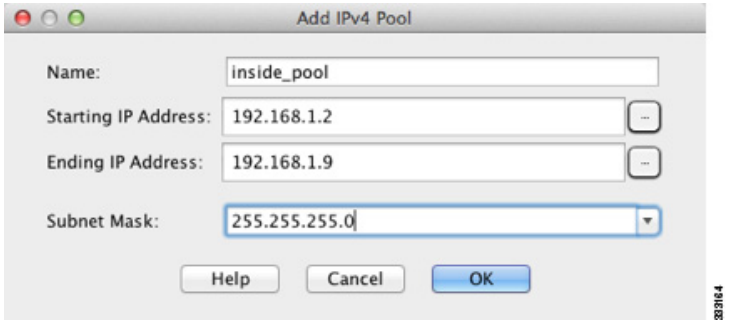
a. In the ASA Cluster area, create a cluster IP pool by clicking the **...** button next to the IP Address Pool field. The valid range shown is determined by the Main IP address you set on the General tab.



b. Click **Add**.

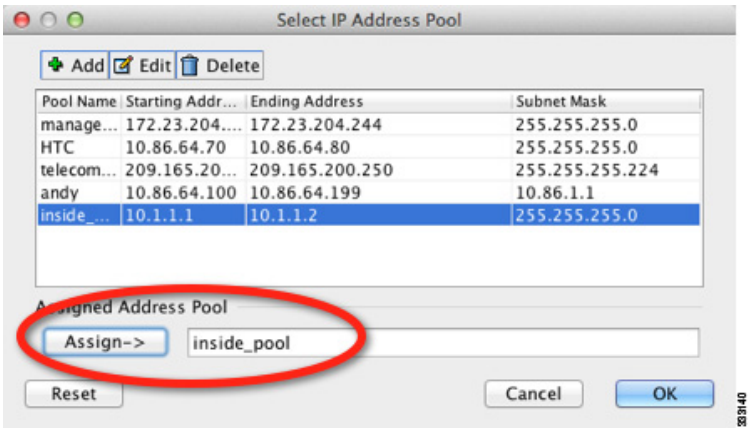


c. Configure a range of addresses that does not include the Main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.



d. Click **OK** to create the new pool.

e. Select the new pool you created, and click **Assign**, and then click **OK**.

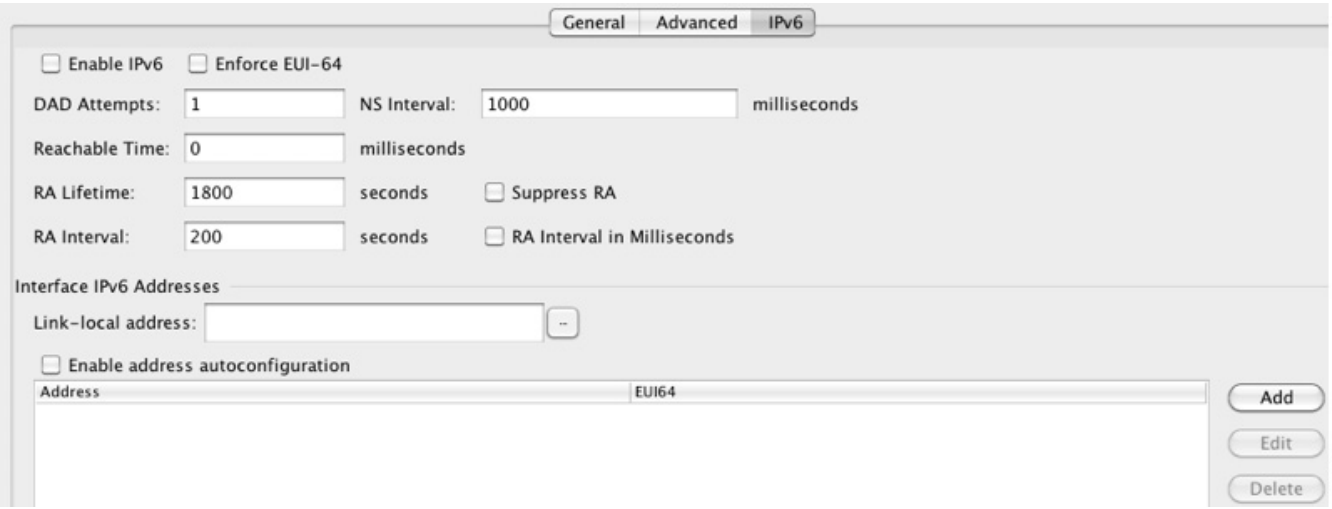


The pool name appears in the IP Address Pool field.

f. (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.

g. For other optional parameters on the Advanced tab, see the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13 and the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 11-25.

Step 4 To configure an IPv6 address, click the **IPv6** tab.



a. Check the **Enable IPv6** checkbox.

b. In the Interface IPv6 Addresses area, click **Add**.

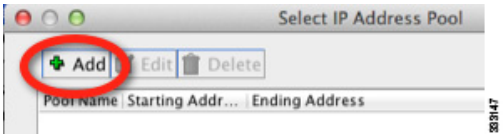
   The Enable address autoconfiguration option is not supported.

   The Add IPv6 Address for Interface dialog box appears.

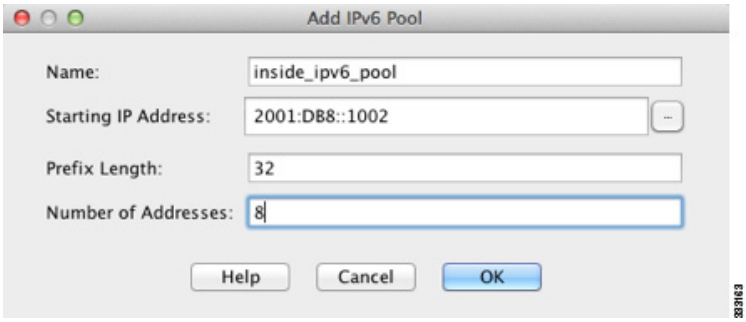
c. In the Address/Prefix Length field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48.



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 13-15.

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 10-43.
- For Individual interface mode, join the cluster. See the "Adding or Joining an ASA Cluster" section on page 10-50.

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 10-35.
- 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 14-8.

Guidelines

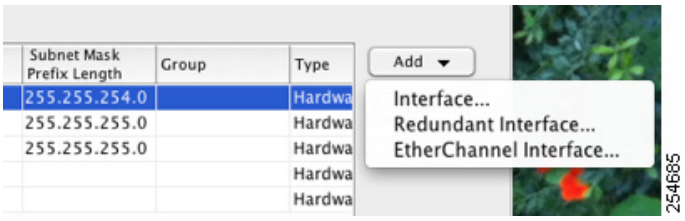
- *Do not* specify the maximum and minimum links in the EtherChannel—We recommend that you do not specify the maximum and minimum links in the EtherChannel on either the ASA or the switch. If you need to use them, note the following:
  - The maximum links set on the ASA is the total number of active 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.
- *Do not* change the load-balancing algorithm from the default. On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Nexus OS and IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the ASAs in a cluster.
- When using Spanned EtherChannels, the port-channel interface will not come up until clustering is fully enabled (see the "Adding or Joining an ASA Cluster" section on page 10-50). This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.
- For detailed EtherChannel guidelines, limitations, and prerequisites, see the "Configuring an EtherChannel" section on page 11-30.
- See also the "EtherChannel Guidelines" section on page 11-13.

Detailed Steps

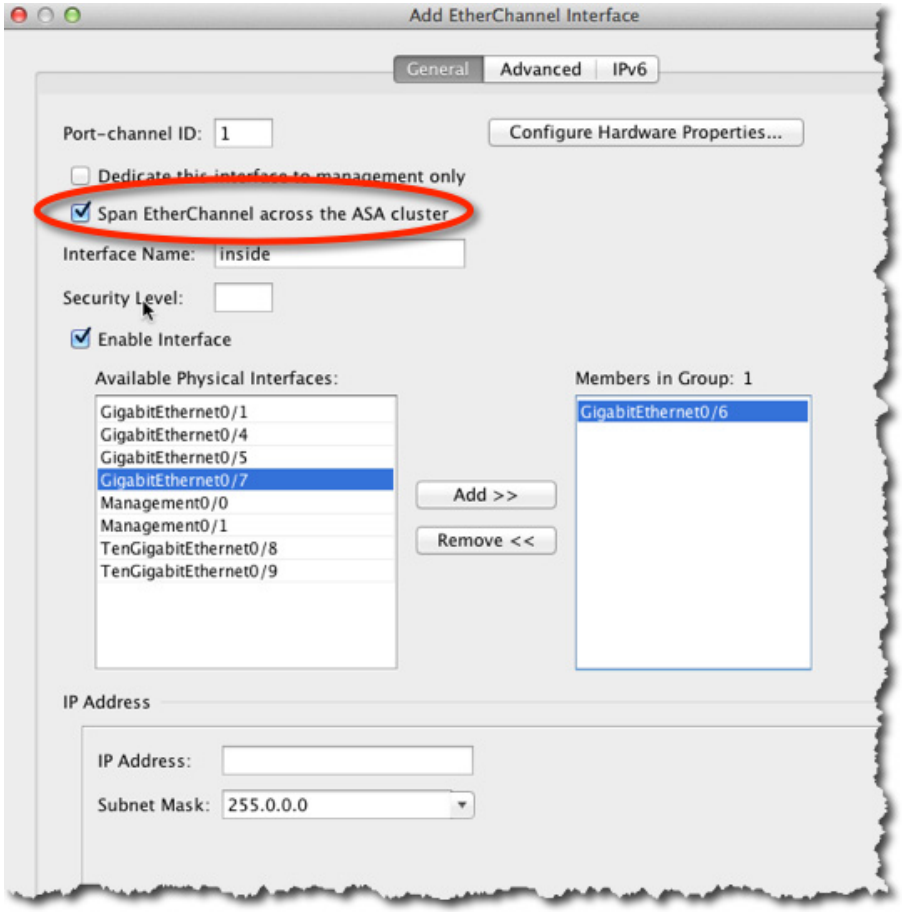
Step 1 Depending on your context mode:

- 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.

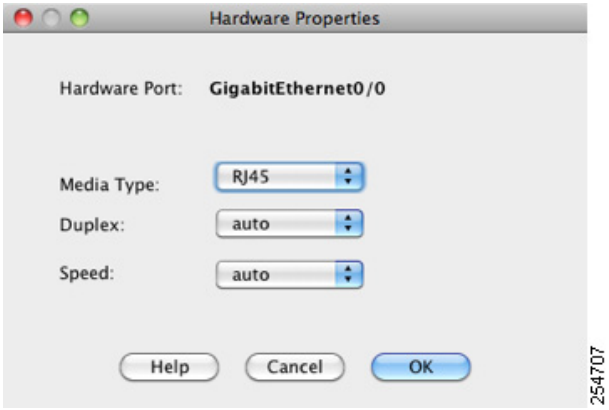


Step 3 Enable the following:

- 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.

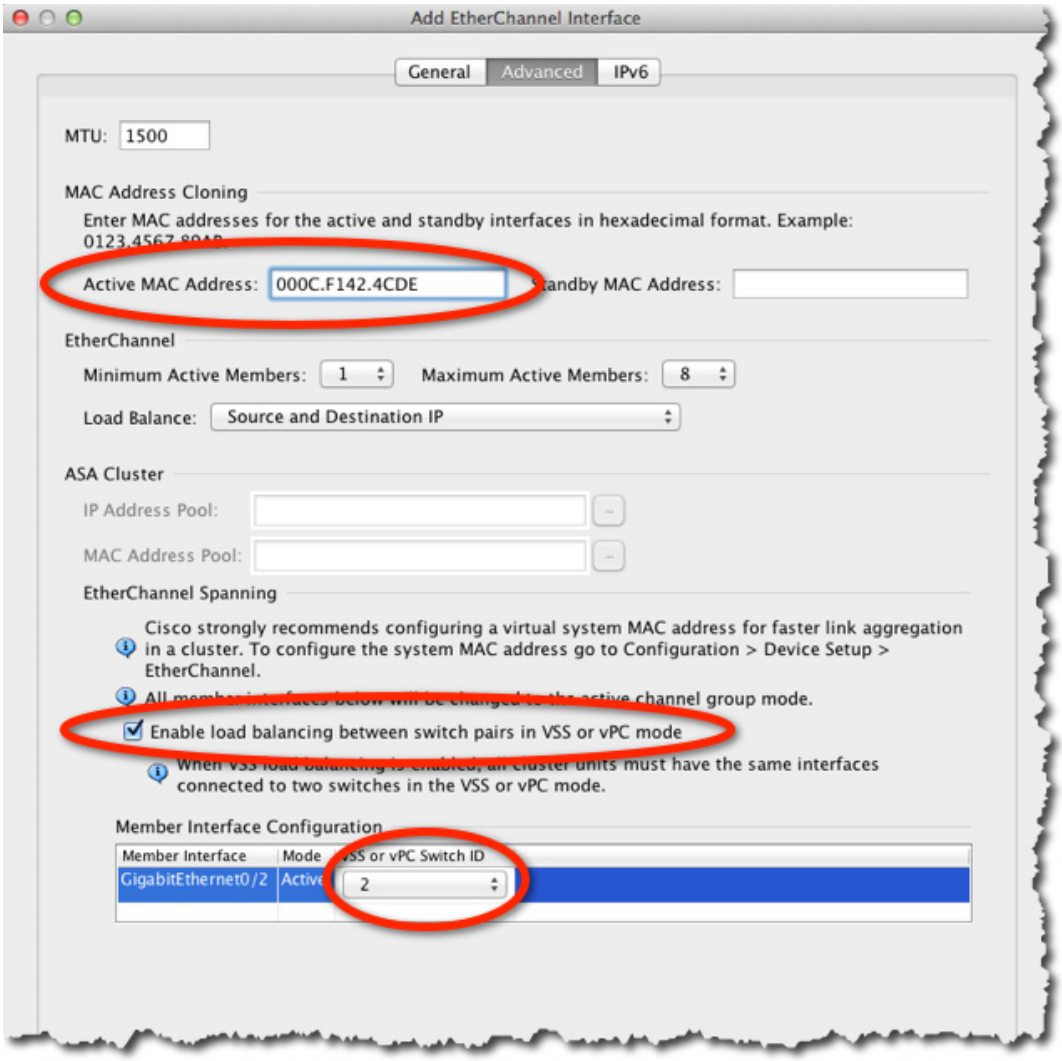
The rest of the fields on this screen are described later in this procedure.

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 11-25.



Click **OK** to accept the Hardware Properties changes.

Step 5 (Optional) To configure the MAC address and optional parameters, click the **Advanced** tab.



a. 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. Note that you must manually configure the MAC address for non-shared interfaces. 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 11-30. For information about the MTU, see the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13.

   We recommend that you do not set the Minimum Active Members and the Maximum Active Members. See the "Guidelines" section on page 10-44 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 11-36. 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 8-20.

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 13-1 for more information.

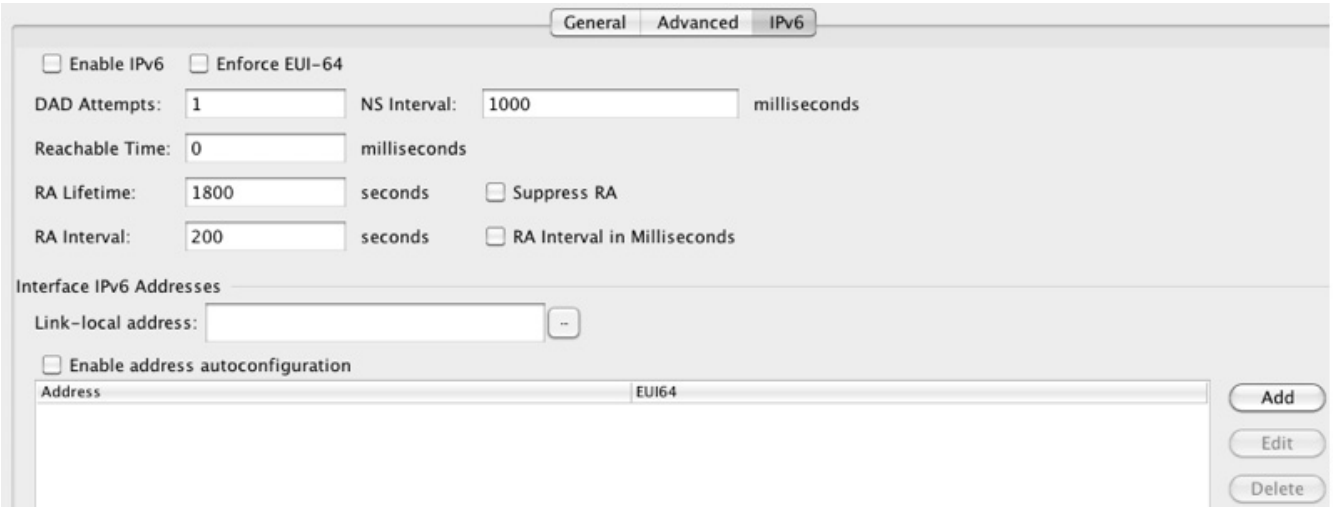
Step 12 (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

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.

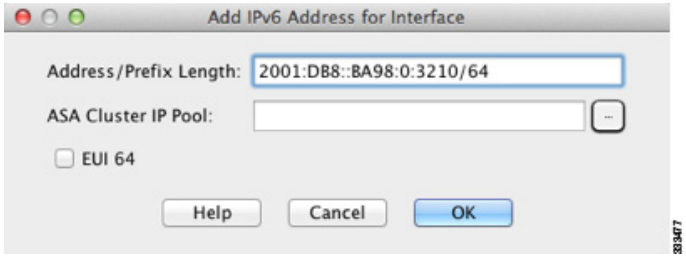


b. In the Interface IPv6 Addresses area, click **Add**.

   The Add IPv6 Address for Interface dialog box appears.

   Note: The Enable address autoconfiguration option is not supported.

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.



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 13-15.

Step 15 Click **OK** to return to the Interfaces screen.

Step 16 Click **Apply**.

What to Do Next

Join the cluster. See the "Adding or Joining an ASA Cluster" section on page 10-50.

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, see the "Configuring ASA Cluster Parameters" section on page 10-59.

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 10-59.

Prerequisites

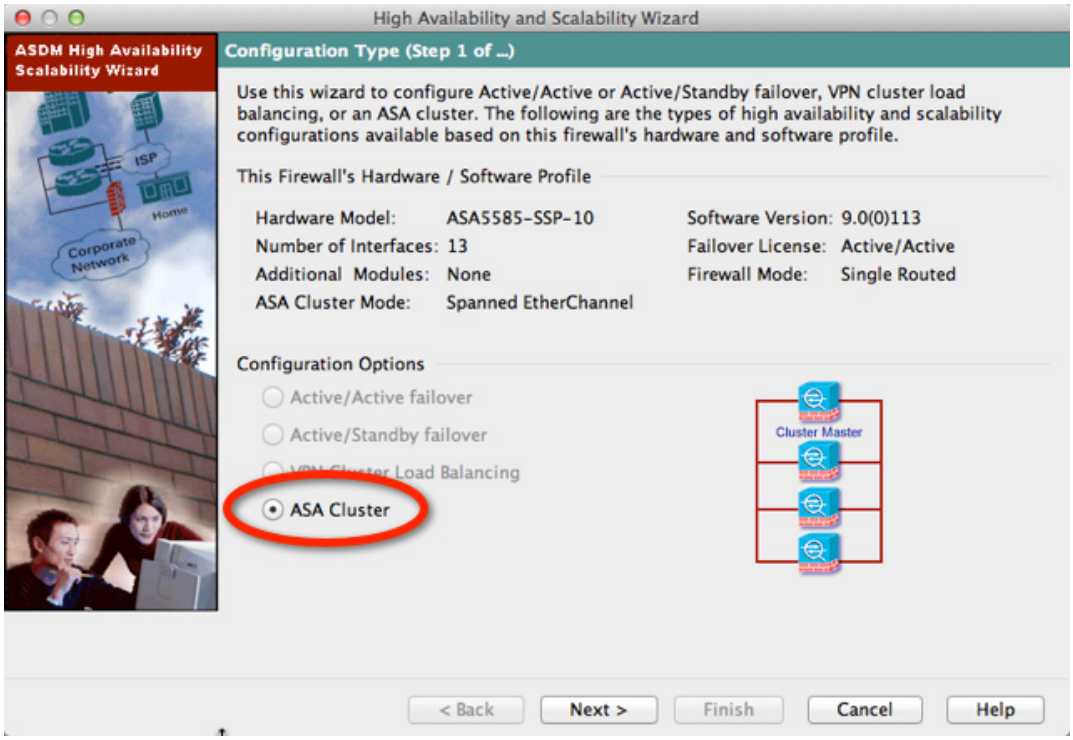
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.
- We recommend enabling jumbo frame reservation for use with the cluster control link. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39.
- The interfaces you intend to use for the cluster control link interface must be in an up state on the connected switch.
- When you add a unit to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.
- We suggest setting the cluster control link MTU to 1600 bytes or greater, which requires you to enable jumbo frame reservation *on each unit* before continuing with this procedure. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39. Jumbo frame reservation requires a reload of the ASA.

Detailed Steps 1—Starting the Wizard

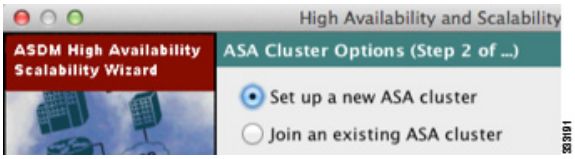
Perform the following steps to start the High Availability and Scalability wizard.

Step 1 Choose **Wizards > High Availability and Scalability Wizard**.

Step 2 Click **ASA Cluster**, and then 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.

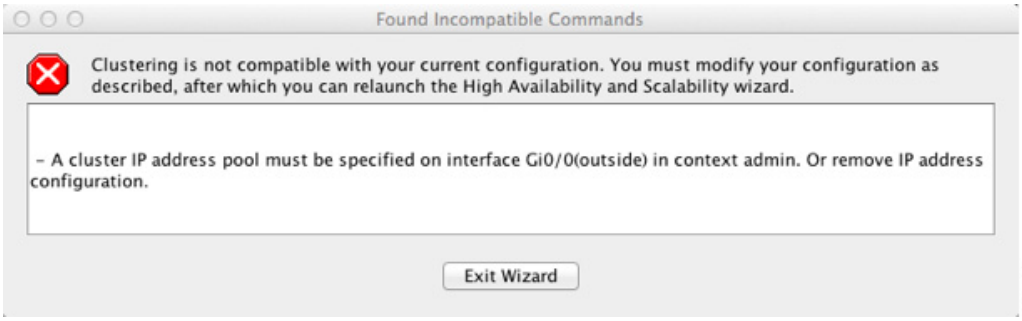
On a master ASA in an existing cluster, the second radio button is labeled **Add another member to the cluster**.

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 10-35), 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.

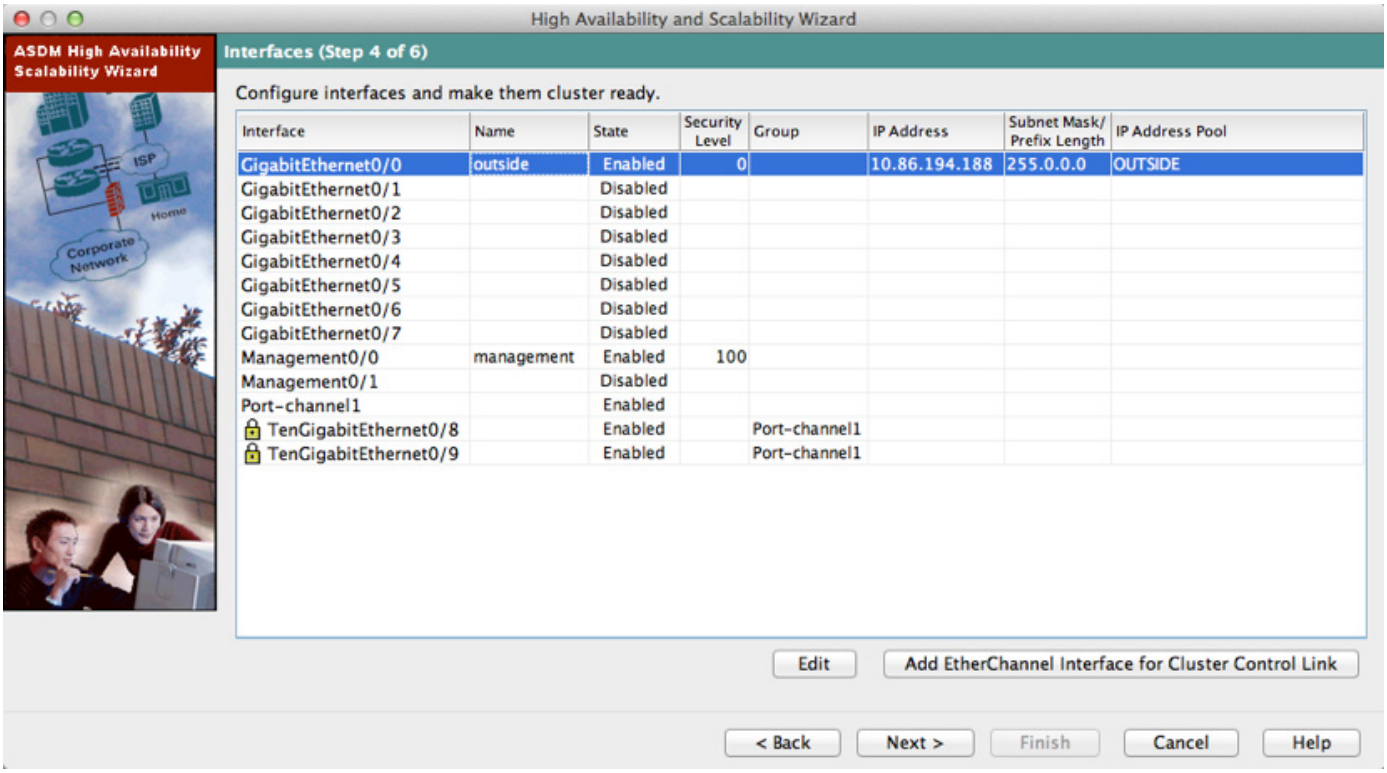
In multiple context mode, if any context interfaces are not compatible with clustering, you see the following error, and are prompted to exit the wizard:



See the "(Recommended; Required in Multiple Context Mode) Configuring Interfaces on the Master Unit" section on page 10-38 to fix your interface configuration.

Detailed Steps 2—Configuring Interfaces

Step 1 Use this screen to configure the cluster control link. In single context mode, you can also use this screen to configure basic interface parameters. In multiple context mode, this screen only lets you configure hardware properties for interfaces.



Step 2 For the cluster control link, you can use a single interface, or add an EtherChannel from this dialog box.

To use a single interface:

a. Select the interface and click **Edit**.

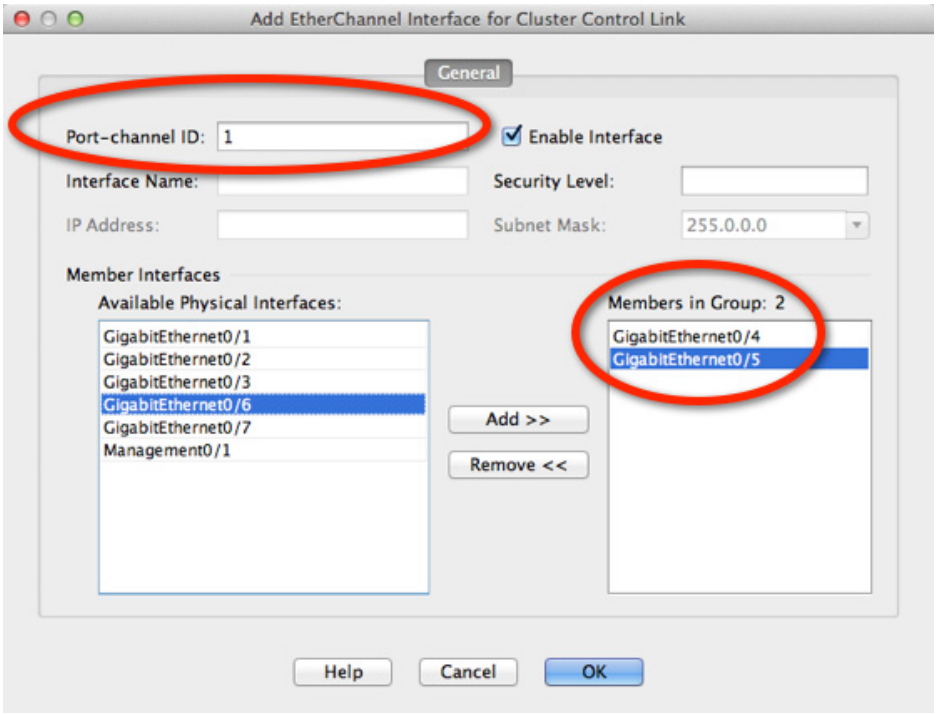
b. Check the **Enable Interface** checkbox, and click **OK**.

Make sure there is no name configured for the interface.

To add an EtherChannel for the cluster control link:

a. Click **Add EtherChannel for Cluster Control Link**.

The Add EtherChannel Interface for Cluster Control Link dialog box appears.



a. Specify the Port-channel ID.

b. Add member interfaces to the Members in Group window.

c. Click **OK** to return to the Interfaces dialog box.

d. Enable each EtherChannel member interface by selecting the interface and clicking **Edit**.

e. Check the **Enable Interface** checkbox, and click **OK**.

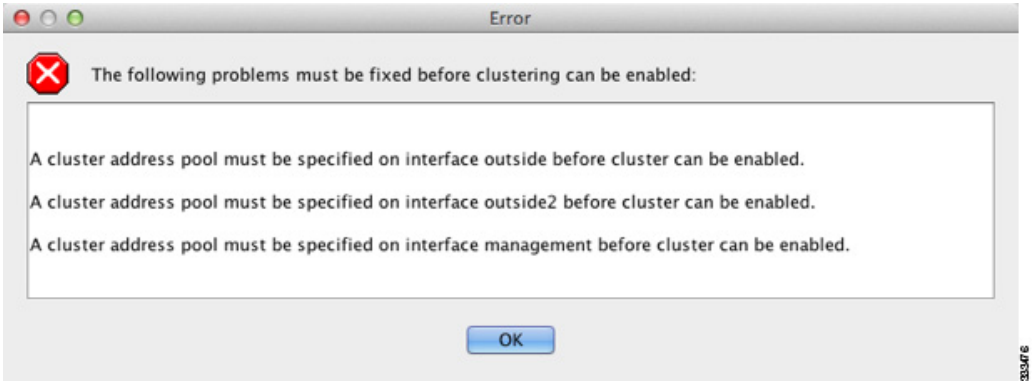
f. Repeat for other members.

Step 3 Make any other 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 Spanned EtherChannels" section on page 10-43
- "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-38

Step 4 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:

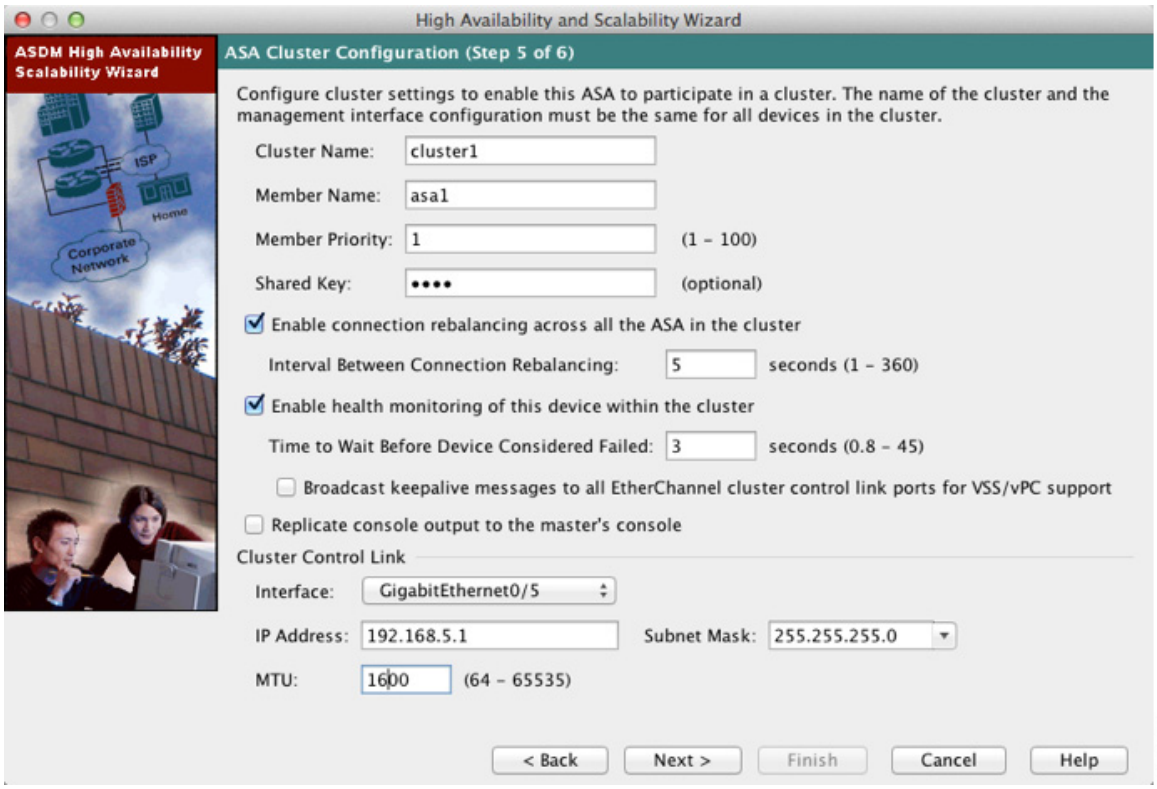


You are returned to the Interfaces screen to fix the interface configuration.

If your interface configuration passes the error checker, the ASA Cluster Configuration screen appears.

Detailed Steps 3—Configuring Bootstrap Settings

Step 1 Configure the following bootstrap settings:



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.

Note Do not configure connection rebalancing for inter-site topologies; you do not want connections rebalanced to cluster members at a different site.

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. To determine unit health, the ASA cluster units send keepalive messages on the cluster control link to other units. If a unit does not receive any keepalive messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead. The interface health check monitors for link failures. 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. If a unit does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. For details, see Interface monitoring, page 10-9.

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.

- Time to Wait Before Device Considered Failed—This value determines the amount of time between unit keepalive status messages, between .8 and 45 seconds; The default is 3 seconds. Note that the holdtime value only affects the *unit* health check; for interface health, the ASA uses the interface status (up or down).
- (Optional) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

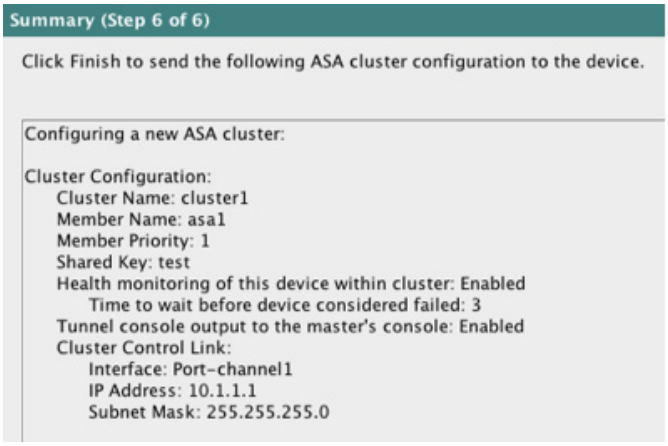
g. (Optional) **Replicate console output to the master's console**—Enables console replication from slave units to the master unit. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, slave units 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. If you already enabled jumbo frame reservation, we suggest setting the MTU to 1600 bytes or greater. If you want to use jumbo frames and have not pre-enabled jumbo frame reservation, you should quit the wizard, enable jumbo frames, and then restart this procedure. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39.

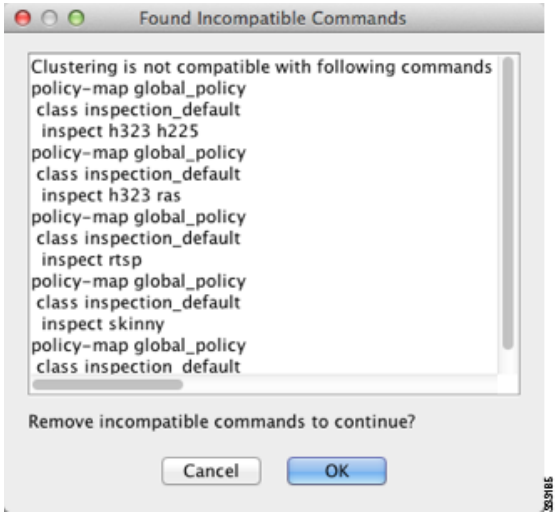
Step 2 Click **Next**.

The wizard shows the cluster configuration.

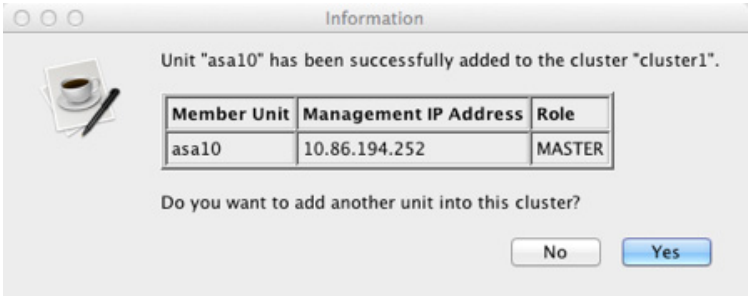


Step 3 Click **Finish**.

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.

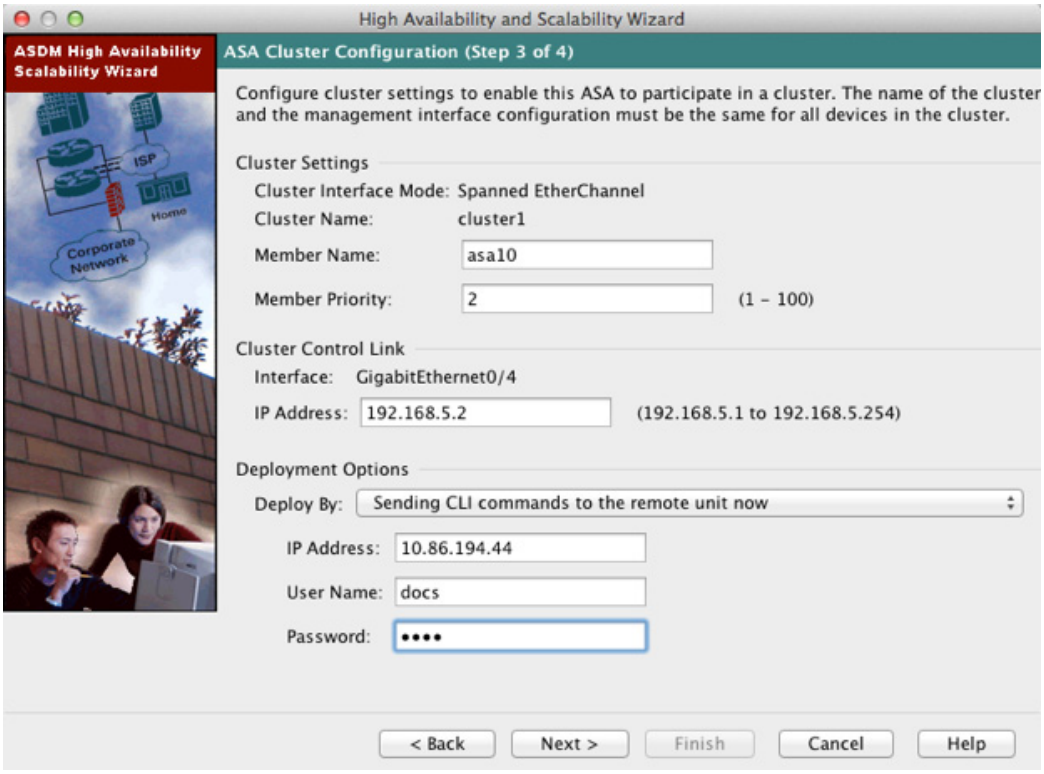


Note In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact 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**.

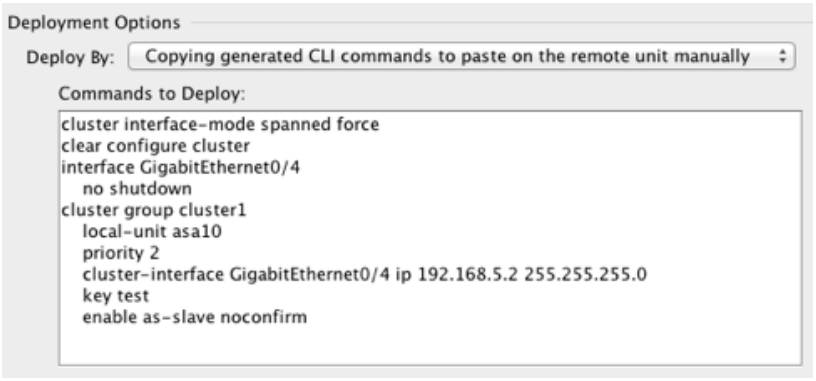
If you are re-running the wizard from the master, you can add slave units by choosing the **Add another member to the cluster** option when you first start the wizard.



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 10-20.

Managing ASA Cluster Members

- Configuring ASA Cluster Parameters, page 10-59
- Adding a New Slave from the Master Unit, page 10-62
- Becoming an Inactive Member, page 10-64
- Inactivating a Slave Member from the Master Unit, page 10-65
- Leaving the Cluster, page 10-66
- Changing the Master Unit, page 10-67
- Executing a Command Cluster-Wide, page 10-69

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

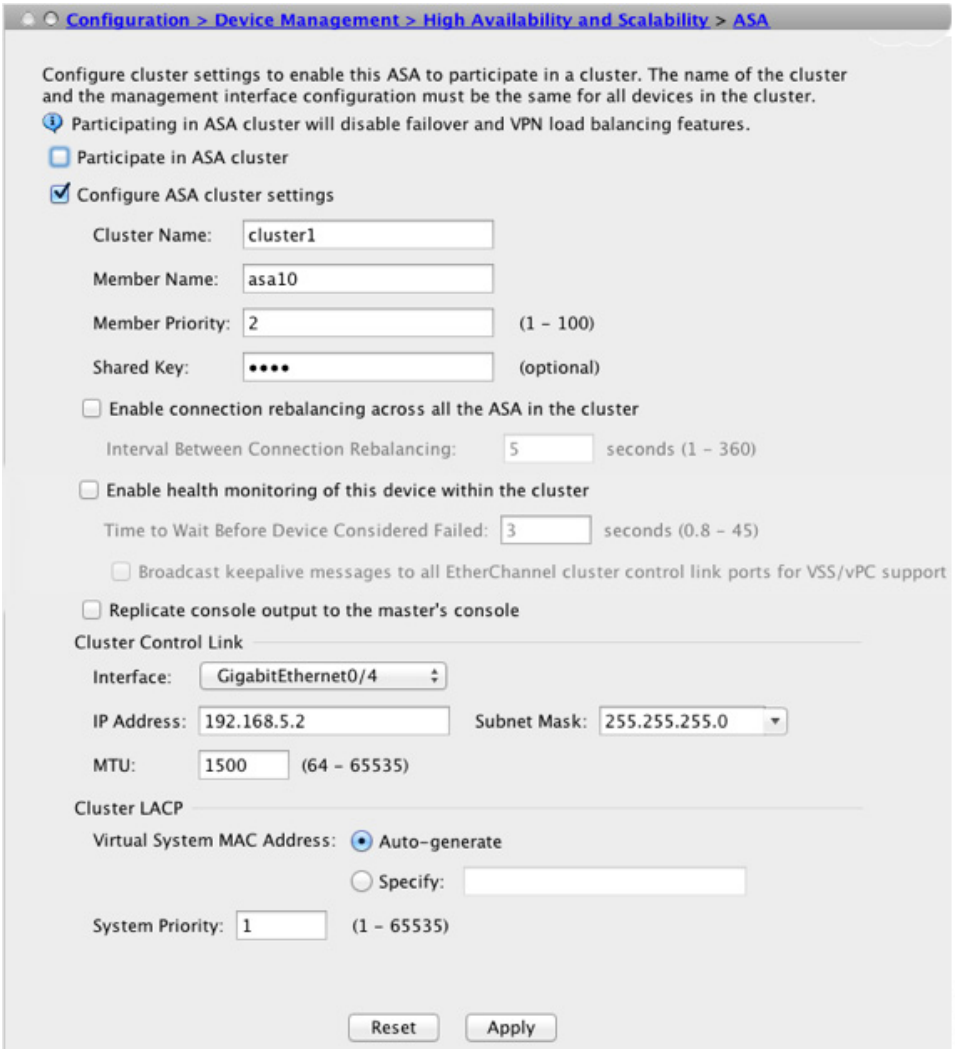
- Pre-configure the cluster control link interfaces on each unit before joining the cluster. For a single interface, you must enable it; do not configure any other settings. For an EtherChannel interface, enable it and set the EtherChannel mode to On.
- For multiple context mode, perform this procedure 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.

Note After you enable clustering, do not uncheck the **Configure ASA cluster settings** checkbox without understanding the consequences. This action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

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. To determine unit health, the ASA cluster units send keepalive messages on the cluster control link to other units. If a unit does not receive any keepalive messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead. Interface status messages detect link failure. 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. If a unit does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. For details, see Interface monitoring, page 10-9.



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.

   – (Optional) Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

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 setting 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 11-39.

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 10-59. However, after you enable clustering, your ASDM session will be disconnected, and you will have to reconnect.

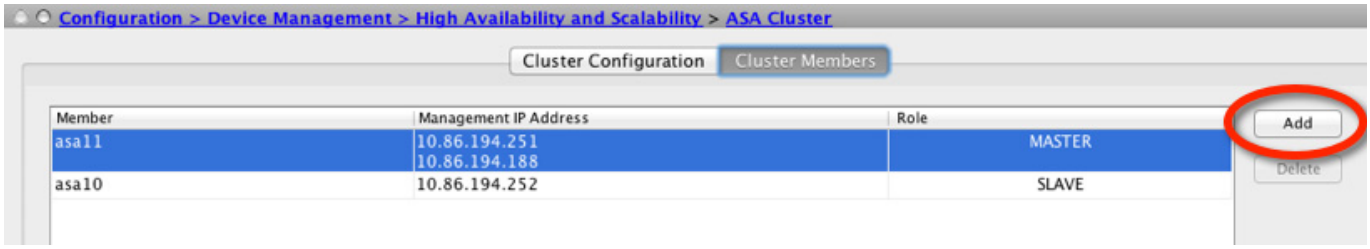
Prerequisites

- 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.

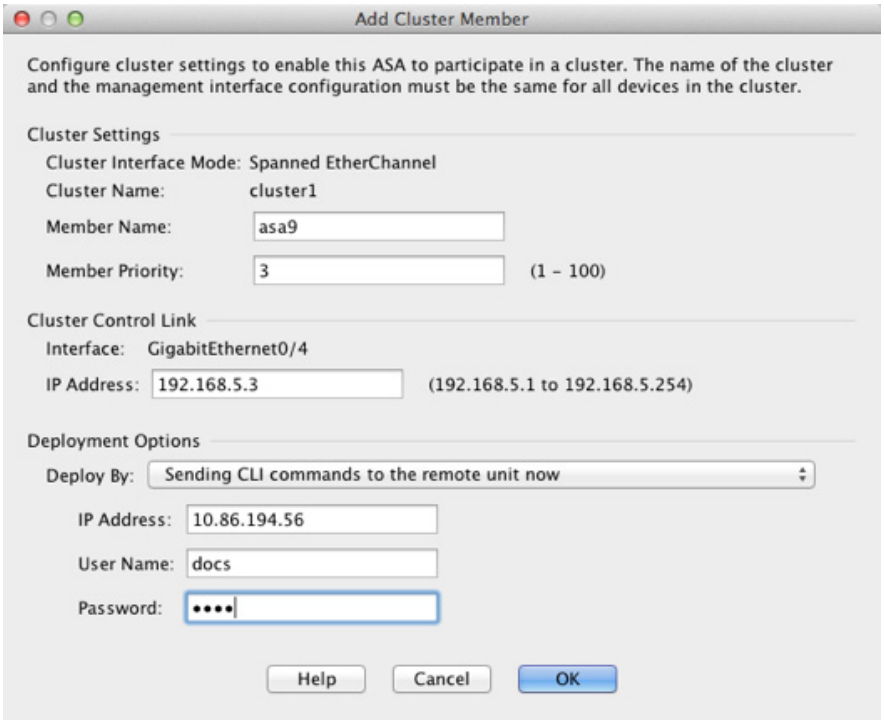
Detailed Steps

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.

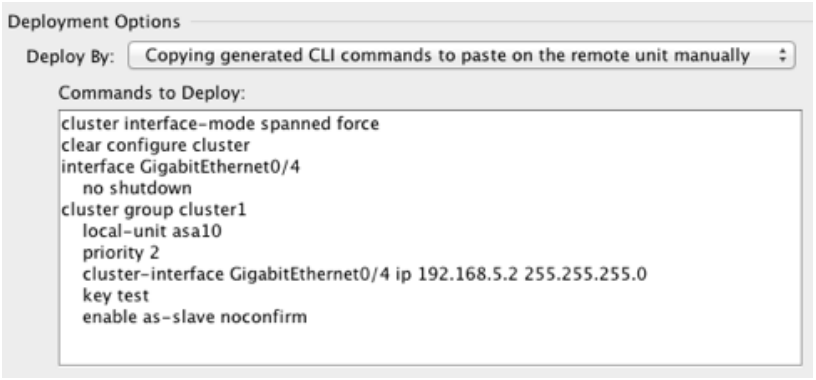
b. Member Priority—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.

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 an ASA becomes inactive (either manually or through a health check failure), 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 10-66. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster, the management interface is not accessible (because it then uses the Main IP address, which is the same 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.

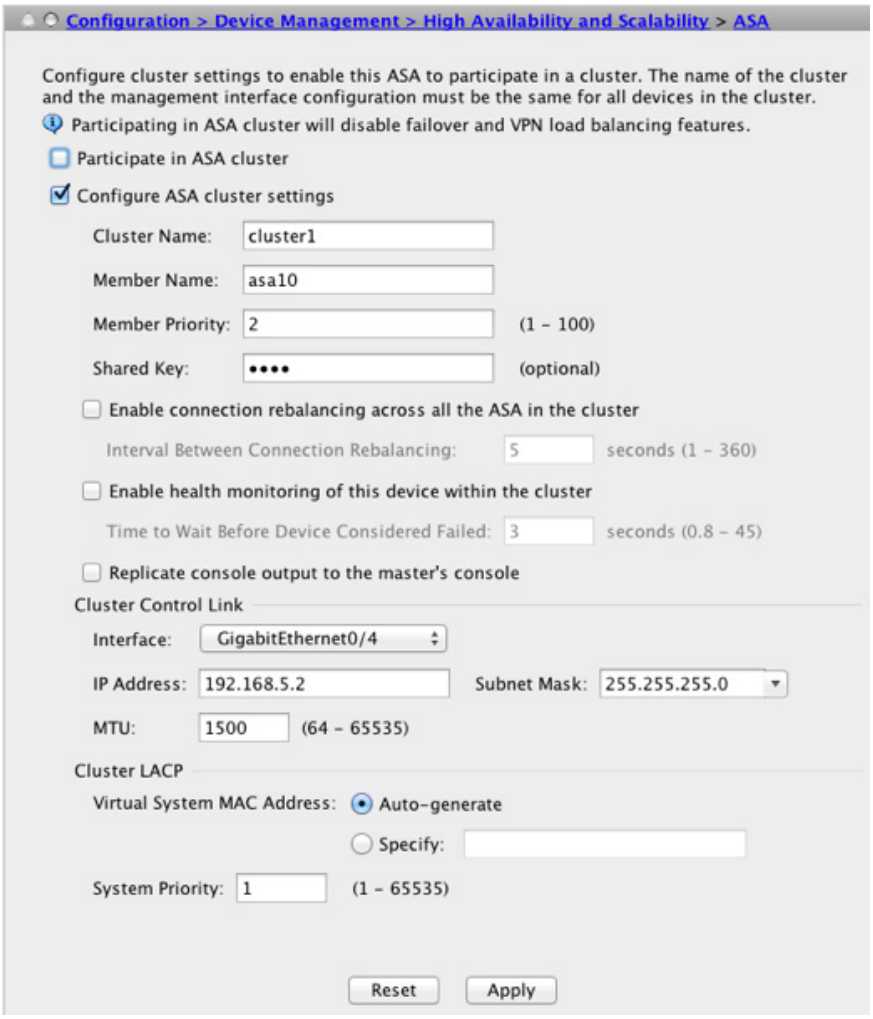
Detailed Steps

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.

Note Do not uncheck the **Configure ASA cluster settings** checkbox; this action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.



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.

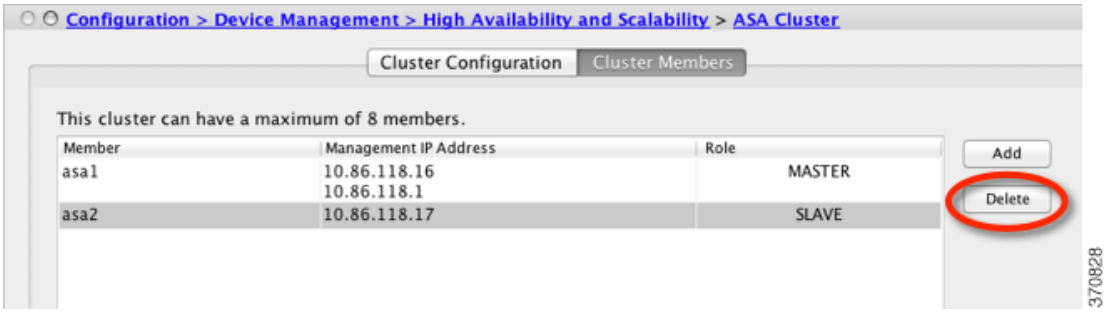
Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the unit altogether from the cluster. See the "Leaving the Cluster" section on page 10-66. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster, the management interface is not accessible (because it then uses the Main IP address, which is the same 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. 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 to avoid IP address conflicts.

Prerequisites

You must use the console port; when you remove the cluster configuration, all interfaces are shut down, including the management interface and cluster control link.

Detailed Steps

	Command	Purpose

Step 1	For a slave unit:
	cluster group *cluster_name*
	no enable
	
	Example:
	`ciscoasa(config)# cluster group cluster1`
	`ciscoasa(cfg-cluster)# no enable`
		Disables clustering. You cannot make configuration changes while clustering is enabled on a slave unit.

Step 2	**clear configure cluster** **Example:** `ciscoasa(config)# clear configure cluster`	Clears the cluster configuration. The ASA shuts down all interfaces including the management interface and cluster control link.
Step 3	**no cluster interface-mode** **Example:** `ciscoasa(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:
	`ciscoasa(config)# copy backup_cluster.cfg running-config`
	
	`Source filename [backup_cluster.cfg]?`
	
	`Destination filename [running-config]?`
	`ciscoasa(config)#`
		Copies the backup configuration to the running configuration.

Step 5	**write memory** **Example:** `ciscoasa(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." Be sure to change the interface IP addresses, and restore the correct hostname, for example.	

Changing the Master Unit

Caution The best method to change the master unit is to disable clustering on the master unit (see the "Becoming an Inactive Member" section on page 10-64), waiting for a new master election, and then re-enabling clustering. If you must specify the exact unit you want to become the master, use the procedure in this section. Note, however, that for centralized features, if you force a master unit change using this procedure, 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 10-21 for a list of centralized features.

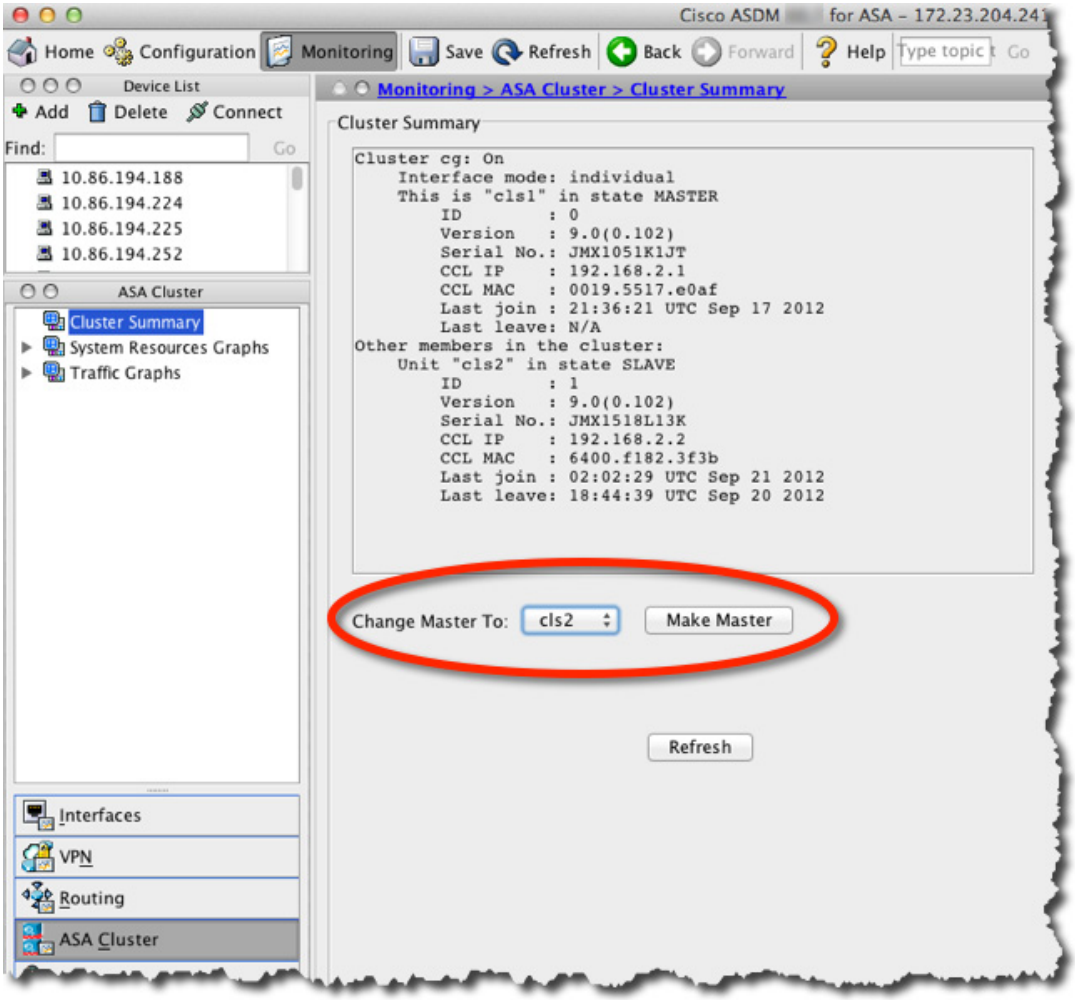
To change the master unit, perform the following steps.

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.

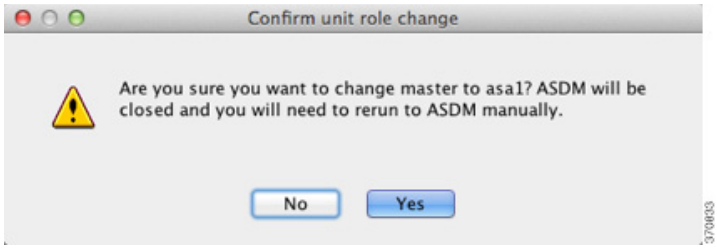
Detailed Steps

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**.

Step 3 You are prompted to confirm the master unit change. Click **Yes**.



Step 4 Quit ASDM, and reconnect using the Main cluster IP address.

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:** `ciscoasa# 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 10-70
- Monitoring Screens, page 10-70
- Related Features, page 10-72

Cluster Dashboards

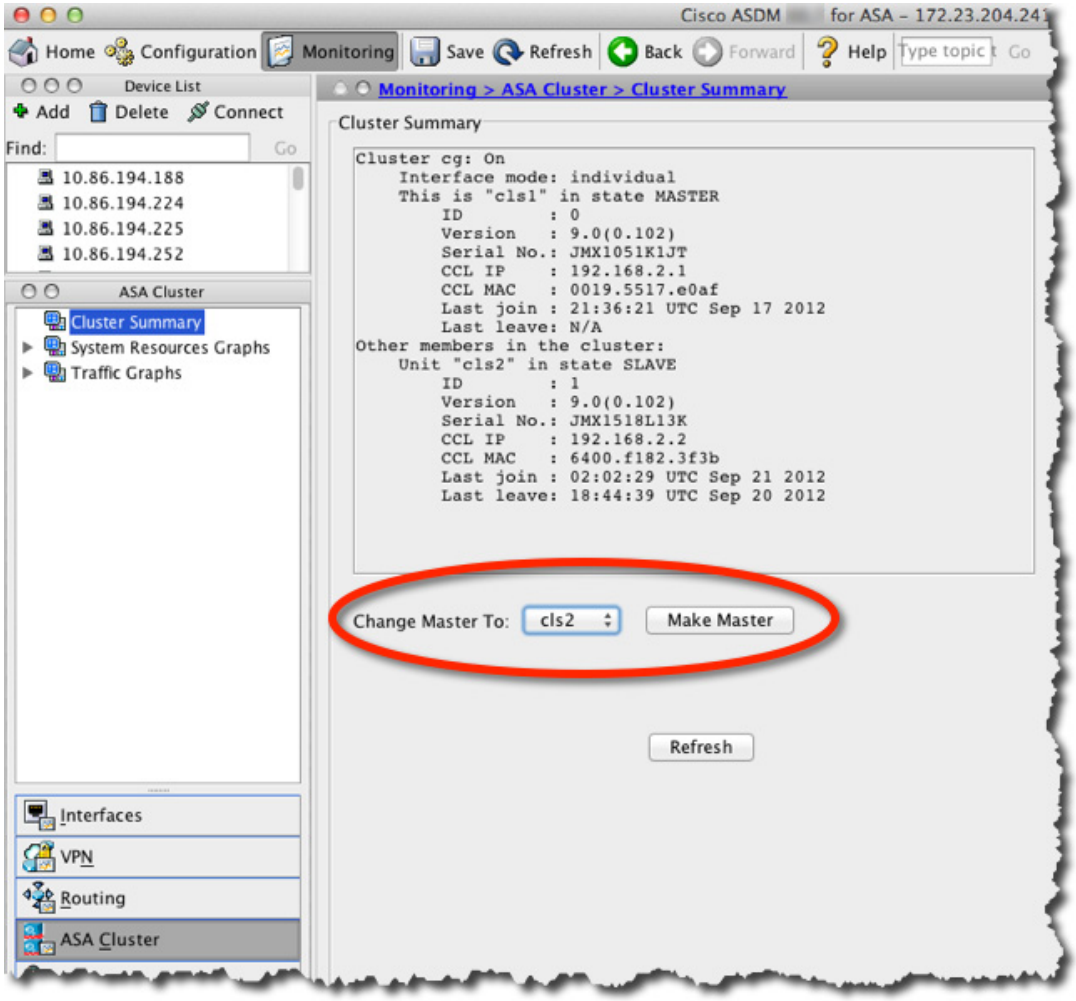
On the Home page on the master unit, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard. For more information, see the "Cluster Dashboard Tab" section on page 4-25 and the "Cluster Firewall Dashboard Tab" section on page 4-26.

Monitoring Screens

- Viewing the Cluster Summary, page 10-70
- Monitoring Cluster Resources, page 10-71
- Monitoring Cluster Traffic, page 10-71
- Monitoring the Cluster Control Link, page 10-72

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 47-1.
Configuration > Device Management > Interfaces > Edit Interface > Advanced	Creates a MAC address pool for an individual interface.
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 45-7.
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 41-22.

Configuration Examples for ASA Clustering

- Sample ASA and Switch Configuration, page 10-72
- Firewall on a Stick, page 10-75
- Traffic Segregation, page 10-77
- Redundant Interface (PBR or ECMP), page 10-80
- Spanned EtherChannel With Backup Links, page 10-82

Sample ASA and Switch Configuration

The following sample configurations connect the following interfaces between the ASA and the switch:

ASA Interface	Switch Interface
GigabitEthernet 0/2	GigabitEthernet 1/0/15
GigabitEthernet 0/3	GigabitEthernet 1/0/16
GigabitEthernet 0/4	GigabitEthernet 1/0/17
GigabitEthernet 0/5	GigabitEthernet 1/0/18

- ASA Configuration, page 10-73
- IOS Switch Configuration, page 10-74

ASA Configuration

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface GigabitEthernet0/0
 channel-group 1 mode on
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
!
cluster group Moya
 local-unit A
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 10
 key emphyri0
 enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface GigabitEthernet0/0
 channel-group 1 mode on
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
!
cluster group Moya
 local-unit B
 cluster-interface Port-channel1 ip 10.0.0.2 255.255.255.0
 priority 11
 key emphyri0
 enable as-slave
```

Master Interface Configuration

```
ip local pool mgmt-pool 10.53.195.231-10.53.195.232

interface GigabitEthernet0/2
 channel-group 10 mode active
 no shutdown
!
interface GigabitEthernet0/3
 channel-group 10 mode active
 no shutdown
!
interface GigabitEthernet0/4
 channel-group 11 mode active
 no shutdown
!
interface GigabitEthernet0/5
 channel-group 11 mode active
 no shutdown
!
interface Management0/0
 management-only
 nameif management
 ip address 10.53.195.230 cluster-pool mgmt-pool
 security-level 100
 no shutdown
!
interface Port-channel10
 port-channel span-cluster
 mac-address aaaa.bbbb.cccc
 nameif inside
 security-level 100
 ip address 209.165.200.225 255.255.255.224
!
interface Port-channel11
 port-channel span-cluster
 mac-address aaaa.dddd.cccc
 nameif outside
 security-level 0
 ip address 209.165.201.1 255.255.255.224
```

IOS Switch Configuration

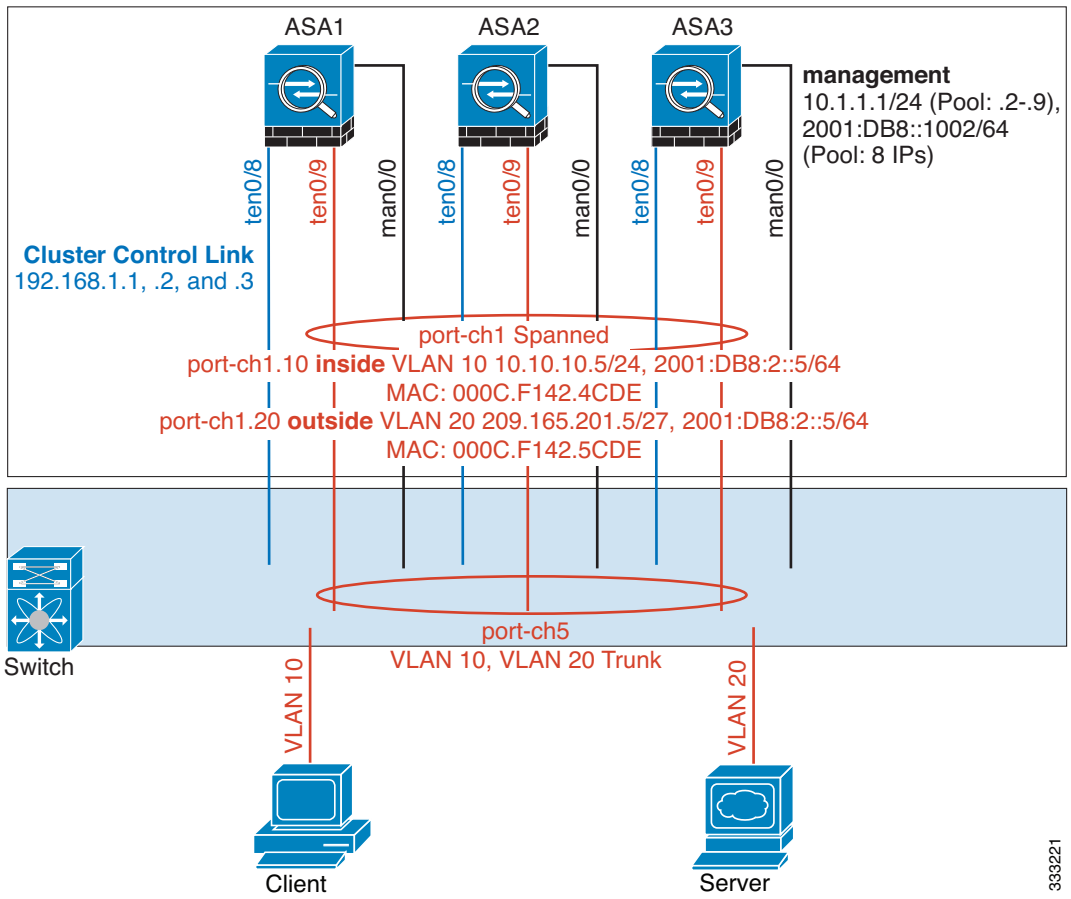
```
interface GigabitEthernet1/0/15
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/16
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/17
 switchport access vlan 401
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active
!
interface GigabitEthernet1/0/18
```

```
 switchport access vlan 401
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active

interface Port-channel10
 switchport access vlan 201
 switchport mode access

interface Port-channel11
 switchport access vlan 401
 switchport mode access
```

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.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/8
    no shutdown
    description CCL

cluster group cluster1
    local-unit asa1
    cluster-interface tengigabitethernet0/8 ip 192.168.1.1 255.255.255.0
    priority 1
    key chuntheunavoidable
    enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/8
    no shutdown
    description CCL

cluster group cluster1
    local-unit asa2
    cluster-interface tengigabitethernet0/8 ip 192.168.1.2 255.255.255.0
    priority 2
    key chuntheunavoidable
    enable as-slave
```

ASA3 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/8
    no shutdown
    description CCL

cluster group cluster1
    local-unit asa3
    cluster-interface tengigabitethernet0/8 ip 192.168.1.3 255.255.255.0
    priority 3
    key chuntheunavoidable
    enable as-slave
```

Master Interface Configuration

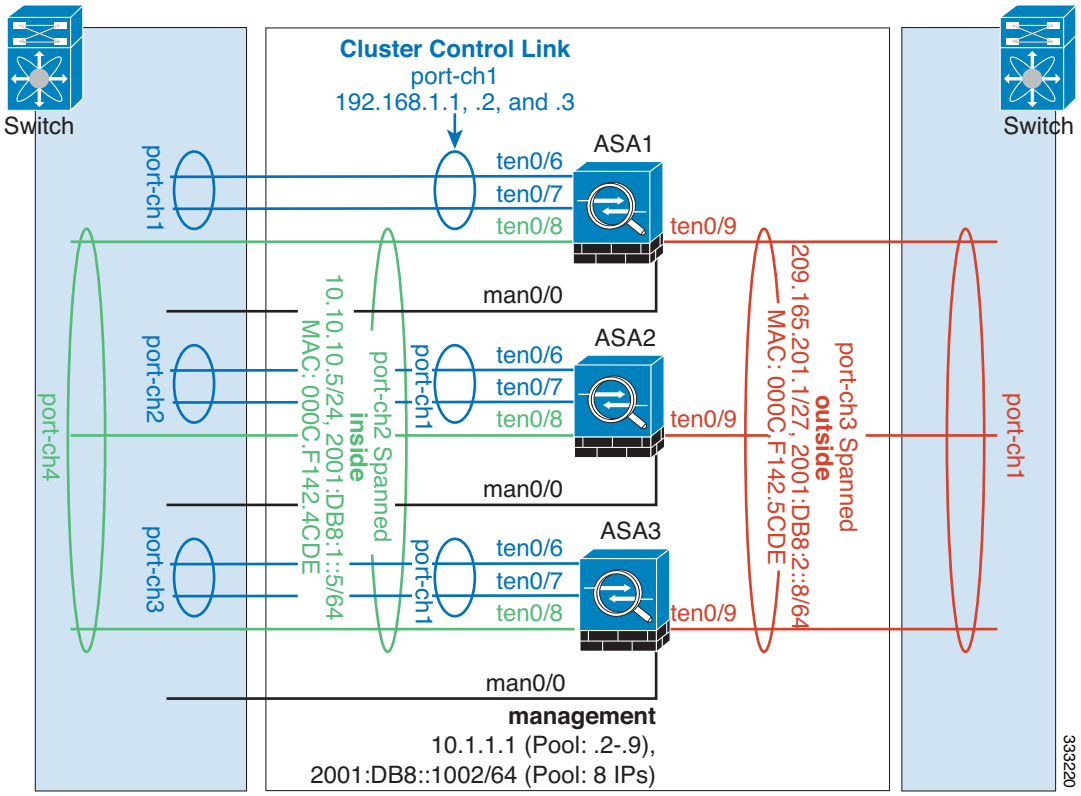
```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

interface management 0/0
    nameif management
    ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
    ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
    security-level 100
    management-only
    no shutdown

interface tengigabitethernet 0/9
    channel-group 2 mode active
    no shutdown
interface port-channel 2
    port-channel span-cluster
interface port-channel 2.10
```

```
    vlan 10
    nameif inside
    ip address 10.10.10.5 255.255.255.0
    ipv6 address 2001:DB8:1::5/64
    mac-address 000C.F142.4CDE
interface port-channel 2.20
    vlan 20
    nameif outside
    ip address 209.165.201.1 255.255.255.224
    ipv6 address 2001:DB8:2::8/64
    mac-address 000C.F142.5CDE
```

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.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
```

```
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa1
    cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
    priority 1
    key chuntheunavoidable
    enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa2
    cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
    priority 2
    key chuntheunavoidable
    enable as-slave
```

ASA3 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa3
    cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
    priority 3
    key chuntheunavoidable
    enable as-slave
```

Master Interface Configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

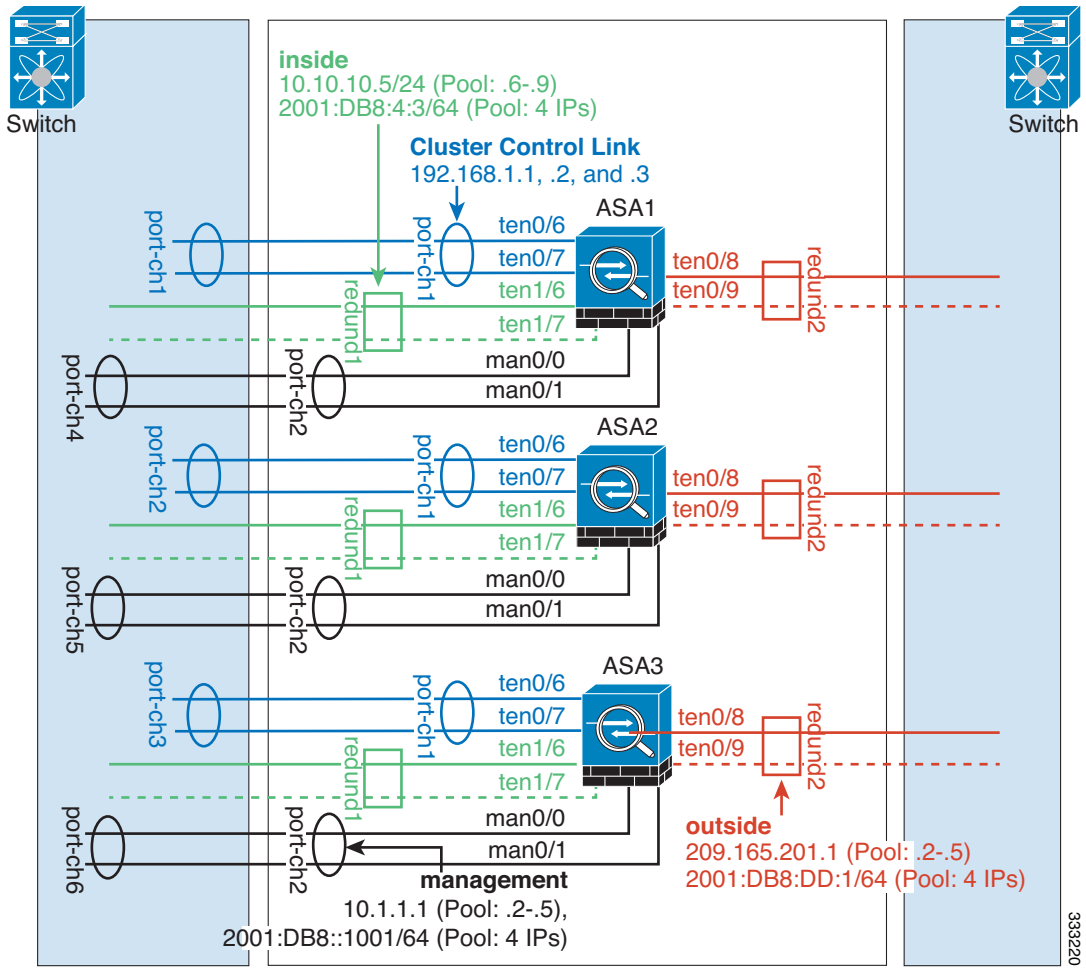
interface management 0/0
    nameif management
    ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
    ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
    security-level 100
    management-only
```

```
    no shutdown

interface tengigabitethernet 0/8
    channel-group 2 mode active
    no shutdown
interface port-channel 2
    port-channel span-cluster
    nameif inside
    ip address 10.10.10.5 255.255.255.0
    ipv6 address 2001:DB8:1::5/64
    mac-address 000C.F142.4CDE

interface tengigabitethernet 0/9
    channel-group 3 mode active
    no shutdown
interface port-channel 3
    port-channel span-cluster
    nameif outside
    ip address 209.165.201.1 255.255.255.224
    ipv6 address 2001:DB8:2::8/64
    mac-address 000C.F142.5CDE
```

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.

Interface Mode on Each Unit

```
cluster interface-mode individual force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
```

```
    local-unit asa1
    cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
    priority 1
    key chuntheunavoidable
    enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa2
    cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
    priority 2
    key chuntheunavoidable
    enable as-slave
```

ASA3 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa3
    cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
    priority 3
    key chuntheunavoidable
    enable as-slave
```

Master Interface Configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.5
ipv6 local pool mgmtipv6 2001:DB8::1002/64 4

interface management 0/0
    channel-group 2 mode active
    no shutdown
interface management 0/1
    channel-group 2 mode active
    no shutdown
interface port-channel 2
    nameif management
    ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
    ipv6 address 2001:DB8::1001/64 cluster-pool mgmtipv6
    security-level 100
    management-only
```

```
ip local pool inside 10.10.10.6-10.10.10.9
ipv6 local pool insideipv6 2001:DB8:4:4/64 4

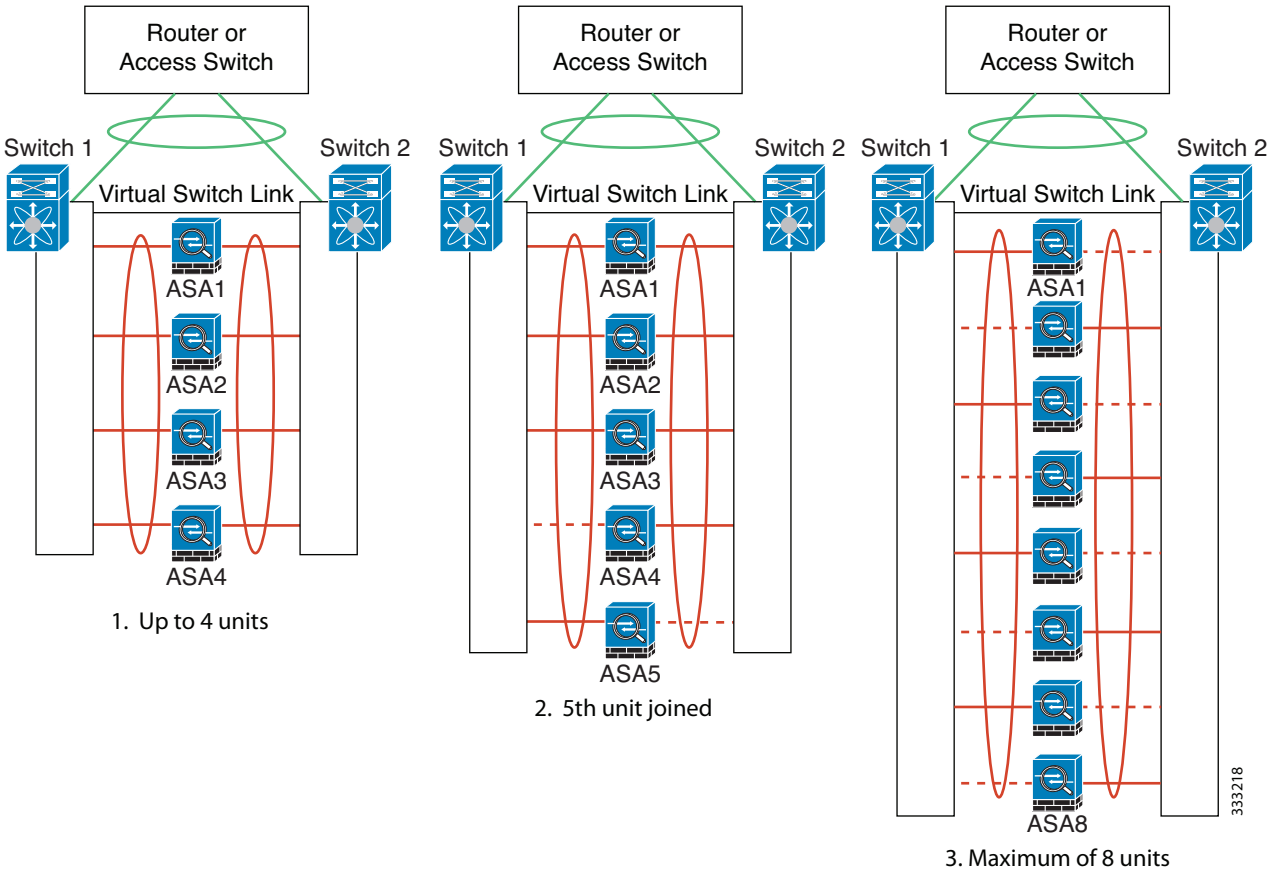
interface redundant 1
member-interface tengigabitethernet 1/6
member-interface tengigabitethernet 1/7
    nameif inside
    ip address 10.10.10.5 255.255.255.0 cluster-pool inside
    ipv6 address 2001:DB8:4:3/64 cluster-pool insideipv6
    security-level 100

ip local pool outside 209.165.201.2-209.165.201.5
ipv6 local pool outsideipv6 2001:DB8:DD:2/64 4

interface redundant 2
member-interface tengigabitethernet 0/8
member-interface tengigabitethernet 0/9
    nameif outside
    ip address 209.165.201.1 255.255.255.224 cluster-pool outside
    ipv6 address 2001:DB8:DD:1/64 cluster-pool outsideipv6
    security-level 0
```

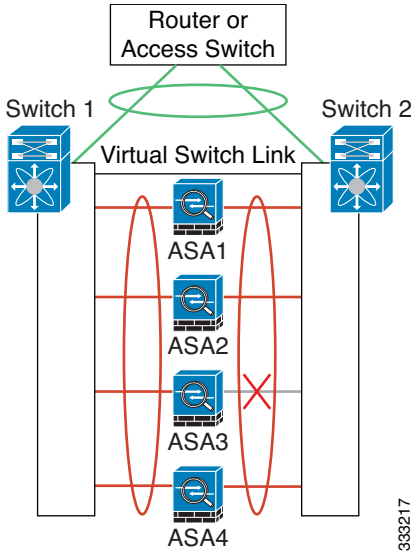
Spanned EtherChannel With Backup Links

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:

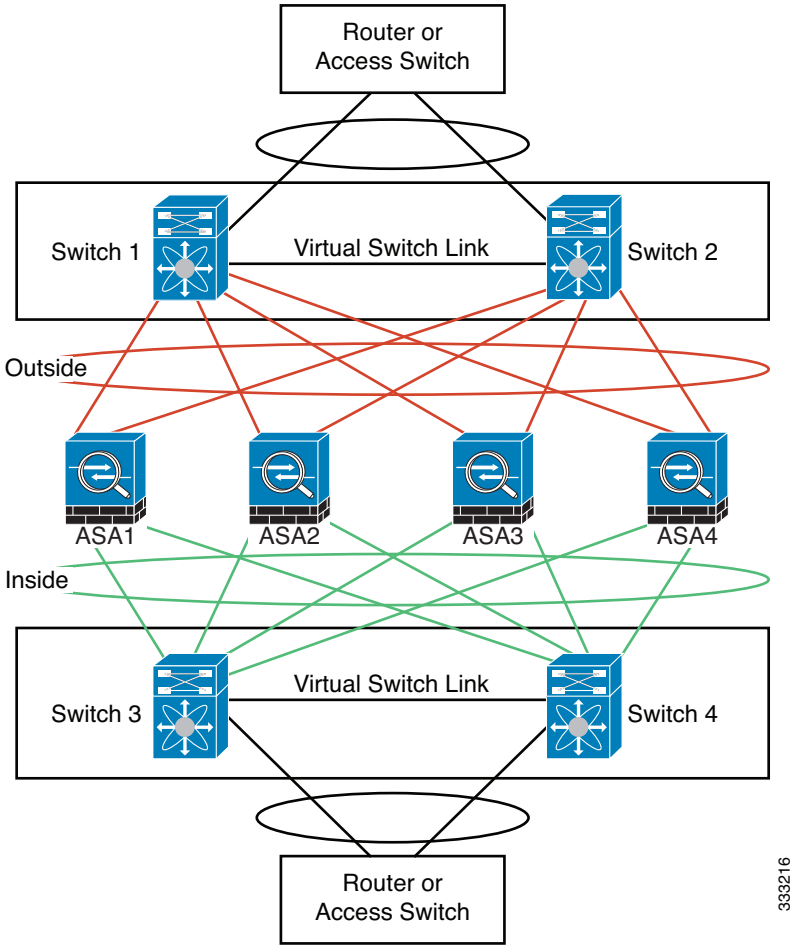


The principle is to first maximize the number of active ports in the channel, and secondly keep the number of active primary ports and the number of active secondary ports in balance. Note that when a 5th unit joins the cluster, traffic is not balanced evenly between all units.

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.



Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/8
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/9
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL
```

```
cluster group cluster1
    local-unit asa1
    cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
    priority 1
    key chuntheunavoidable
    enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/8
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/9
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa2
    cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
    priority 2
    key chuntheunavoidable
    enable as-slave
```

ASA3 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/8
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/9
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa3
    cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
    priority 3
    key chuntheunavoidable
    enable as-slave
```

ASA4 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/6
    channel-group 1 mode on
```

```
    no shutdown
interface tengigabitethernet 0/7
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/8
    channel-group 1 mode on
    no shutdown
interface tengigabitethernet 0/9
    channel-group 1 mode on
    no shutdown
interface port-channel 1
    description CCL

cluster group cluster1
    local-unit asa4
    cluster-interface port-channel1 ip 192.168.1.4 255.255.255.0
    priority 4
    key chuntheunavoidable
    enable as-slave
```

Master Interface Configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9

interface management 0/0
    channel-group 2 mode active
    no shutdown
interface management 0/1
    channel-group 2 mode active
    no shutdown
interface port-channel 2
    nameif management
    ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
    security-level 100
    management-only

interface tengigabitethernet 1/6
    channel-group 3 mode active vss-id 1
    no shutdown
interface tengigabitethernet 1/7
    channel-group 3 mode active vss-id 2
    no shutdown
interface port-channel 3
    port-channel span-cluster vss-load-balance
    nameif inside
    ip address 10.10.10.5 255.255.255.0
    mac-address 000C.F142.4CDE

interface tengigabitethernet 1/8
    channel-group 4 mode active vss-id 1
    no shutdown
interface tengigabitethernet 1/9
    channel-group 4 mode active vss-id 2
    no shutdown
interface port-channel 4
    port-channel span-cluster vss-load-balance
    nameif outside
    ip address 209.165.201.1 255.255.255.224
    mac-address 000C.F142.5CDE
```

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 for the ASA 5580 and 5585-X	9.0(1)	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
ASA 5500-X support for clustering	9.1(4)	The ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X now support 2-unit clusters. Clustering for 2 units is enabled by default in the base license; for the ASA 5512-X, you need the Security Plus license. We did not modify any ASDM screens.

Table 10-3 Feature History for Clustering (continued)

Feature Name	Platform Releases	Feature Information
Improved VSS and vPC support for health check monitoring	9.1(4)	If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, you can now increase stability with health check monitoring. For some switches, such as the Nexus 5000, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable the VSS/vPC health check feature, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster
Support for cluster members at different geographical locations (inter-site); Individual Interface mode only	9.1(4)	You can now place cluster members at different geographical locations when using individual interface mode. We did not modify any ASDM screens.



PART 3

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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

For ASA 5505 configuration, see Chapter 12, "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 11-2
- Licensing Requirements for ASA 5510 and Higher Interfaces, page 11-10
- Guidelines and Limitations, page 11-12
- Default Settings, page 11-14
- Starting Interface Configuration (ASA 5510 and Higher), page 11-14
- Monitoring Interfaces, page 11-40
- Where to Go Next, page 11-43
- Feature History for ASA 5510 and Higher Interfaces, page 11-44

Information About Starting ASA 5510 and Higher Interface Configuration

This section includes the following topics:

- Auto-MDI/MDIX Feature, page 11-2
- Interfaces in Transparent Mode, page 11-2
- Management Interface, page 11-2
- Redundant Interfaces, page 11-5
- EtherChannels, page 11-5
- Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size, page 11-8

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 14-2.

Management Interface

- Management Interface Overview, page 11-2
- Management Slot/Port Interface, page 11-3
- Using Any Interface for Management-Only Traffic, page 11-3
- Management Interface for Transparent Mode, page 11-4
- No Support for Redundant Management Interfaces, page 11-4
- Management 0/0 Interface on the ASA 5512-X through ASA 5555-X, page 11-4

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 45, "Configuring Management Access."

Management *Slot*/*Port* Interface

Table 11-1 shows the Management interfaces per model.-

Table 11-1 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 3-19 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, MTU, and TCP MSS" section on page 13-13 or the "Configuring Multiple Contexts" section on page 8-15). 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 11-5
- Connecting to an EtherChannel on Another Device, page 11-5
- Link Aggregation Control Protocol, page 11-6
- Load Balancing, page 11-7
- EtherChannel MAC Address, page 11-8

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.

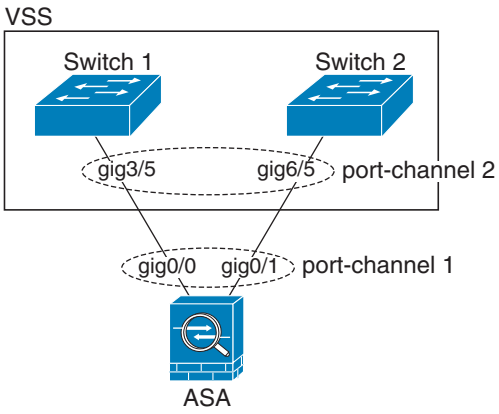
The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The port is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and vlan numbers.

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.

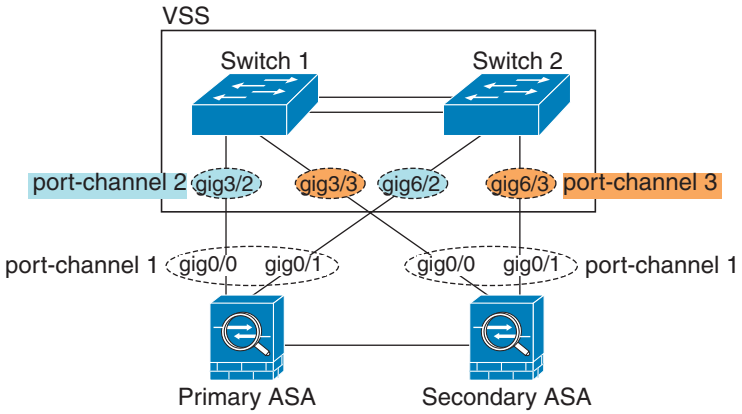
When the switch is part of a Virtual Switching System (VSS), then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch (see Figure 11-1).

Figure 11-1 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 11-1). 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 11-2 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 11-33). 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 11-2 shows the load balancing amounts per interface for each number of active interfaces. The active interfaces in **bold** have even distribution.

Table 11-2 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.

Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size

- MTU Overview, page 11-8
- Default MTU, page 11-8
- Path MTU Discovery, page 11-9
- Setting the MTU and Jumbo Frames, page 11-9
- TCP Maximum Segment Size Overview, page 11-9
- Default TCP MSS, page 11-9
- Setting the TCP MSS for VPN and Non-VPN Traffic, page 11-9

MTU Overview

The maximum transmission unit (MTU) specifies the maximum frame payload size (without Ethernet headers) that the ASA can transmit on a given Ethernet interface. If an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

Note The ASA can receive frames larger than the configured MTU as long as there is room in memory. See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39 to increase memory for larger frames.

Default MTU

The default MTU on the ASA is 1500 bytes. This value does not include the 18 or more bytes for the Ethernet header, CRC, VLAN tagging, and so on.

Path MTU Discovery

The ASA supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

Setting the MTU and Jumbo Frames

See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13. For multiple context mode, set the MTU within each context.

See the "Enabling Jumbo Frame Support (Supported Models)" section on page 11-39. For multiple context mode, set the jumbo frame support in the system execution space.

See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all ASA interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—If your model supports jumbo frames, you can set the MTU up to 9216 bytes.

TCP Maximum Segment Size Overview

The TCP maximum segment size (TCP MSS) is the size of the TCP payload before any TCP headers are added. UDP packets are not affected. The client and the server exchange TCP MSS values during the three-way handshake when establishing the connection.

You can set the TCP MSS on the ASA. If either endpoint of a connection requests a TCP MSS that is larger than the value set on the ASA, the ASA overwrites the TCP MSS in the request packet with the ASA maximum. If the host or server does not request a TCP MSS, then the ASA assumes the RFC 793-default value of 536 bytes, but does not modify the packet. You can also configure the minimum TCP MSS; if a host or server requests a very small TCP MSS, the ASA can adjust the value up. By default, the minimum TCP MSS is not enabled.

For example, you configure the default MTU of 1500 bytes. A host requests an MSS of 1700. If the ASA maximum TCP MSS is 1380, then the ASA changes the MSS value in the TCP request packet to 1380. The server then sends 1380-byte packets.

Default TCP MSS

By default, the maximum TCP MSS on the ASA is 1380 bytes. This default accommodates VPN connections where the headers can add up to 120 bytes; this value fits within the default MTU of 1500 bytes.

Setting the TCP MSS for VPN and Non-VPN Traffic

See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13. For multiple context mode, set the TCP MSS within each context.

See the following guidelines:

- Non-VPN traffic—If you do not use VPN and do not need extra space for headers, then you should disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-VPN packets usually fit this TCP MSS.
- VPN traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to a higher value, then you need to set the TCP MSS to accommodate the new MTU.

Licensing Requirements for ASA 5510 and Higher Interfaces

Model	License Requirement
ASA 5510	VLANs[1]: 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[2]: Base License: 364 Security Plus License: 564
ASA 5520	VLANs[1]: Base License: 150. Interfaces of all types[2]: Base License: 764
ASA 5540	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 964
ASA 5550	VLANs[1]: Base License: 400 Interfaces of all types[2]: Base License: 1764
ASA 5580	VLANs[1]: Base License: 1024 Interfaces of all types[2]: Base License: 4612

Model	License Requirement
ASA 5512-X	VLANs[1]: Base License: 50 Security Plus License: 100 Interfaces of all types[2]: Base License: 716 Security Plus License: 916
ASA 5515-X	VLANs[1]: Base License: 100 Interfaces of all types[2]: Base License: 916
ASA 5525-X	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 1316
ASA 5545-X	VLANs[1]: Base License: 300 Interfaces of all types[2]: Base License: 1716
ASA 5555-X	VLANs[1]: Base License: 500 Interfaces of all types[2]: Base License: 2516
ASA 5585-X	VLANs[1]: Base and Security Plus 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[2]: Base and Security Plus License: 4612

1. For an interface to count against the VLAN limit, you must assign a VLAN to it.
2. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** defined in the configuration counts against this limit.

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 11-14. Then, configure the logical interface parameters in the context execution space according to Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 10-43.
- To configure an individual cluster interface, see the "Configuring Individual Interfaces (Recommended for the Management Interface)" section on page 10-38.

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 11-12.
- For clustering guidelines, see the "Clustering Guidelines" section on page 11-12.

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 11-12.
- For clustering guidelines, see the "Clustering Guidelines" section on page 11-12.

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 3-19.

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 11-15
- Converting In-Use Interfaces to a Redundant or EtherChannel Interface, page 11-16
- Enabling the Physical Interface and Configuring Ethernet Parameters, page 11-25
- Configuring a Redundant Interface, page 11-27

- Configuring an EtherChannel, page 11-30
- Configuring VLAN Subinterfaces and 802.1Q Trunking, page 11-36
- Enabling Jumbo Frame Support (Supported Models), page 11-39

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 11-16.

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 11-25.

Physical interfaces are disabled by default.

Step 3 (Optional) Configure redundant interface pairs. See the "Configuring a Redundant Interface" section on page 11-27.

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 11-30.

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 11-36.

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 11-39.

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 8, "Configuring Multiple Context Mode":

- To assign interfaces to contexts, see the "Configuring a Security Context" section on page 8-20.
- (Optional) To automatically assign unique MAC addresses to context interfaces, see the "Automatically Assigning MAC Addresses to Context Interfaces" section on page 8-24.

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, MTU, and TCP MSS" section on page 13-13.

Step 8 Complete the interface configuration according to Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 11-27 and the "Configuring an EtherChannel" section on page 11-30 fore more information.

- Detailed Steps (Single Mode), page 11-16
- Detailed Steps (Multiple Mode), page 11-21

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 11-30.

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 46-25 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 46-21.

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 11-27 or the "Configuring an EtherChannel" section on page 11-30. 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 8-20.

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

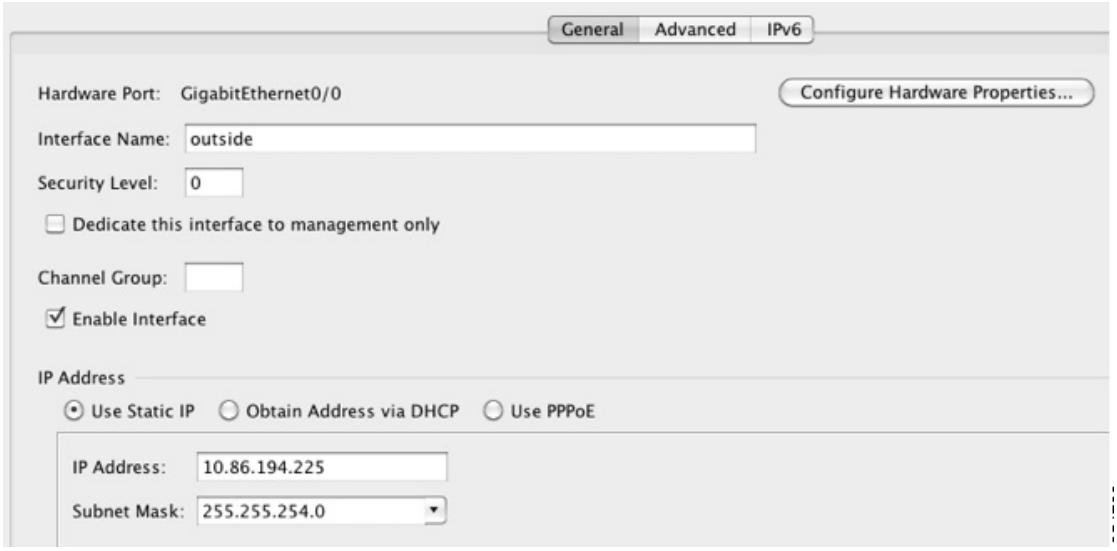
Step 1 Depending on your context mode:

- 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.



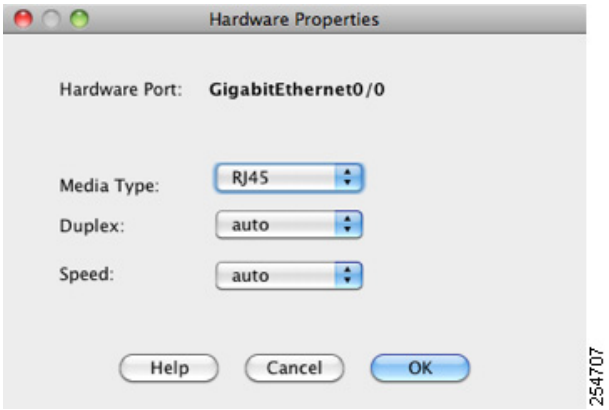
Note In single mode, this procedure only covers a subset of the parameters on the Edit Interface dialog box; to configure other parameters, see Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)." Note that 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 8-15.

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 11-2.

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 11-27.
- Configure an EtherChannel. See the "Configuring an EtherChannel" section on page 11-30.
- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 11-36.

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 8-15.
- For single context mode, complete the interface configuration. See Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 11-27
- Changing the Active Interface, page 11-30

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 11-13.

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.

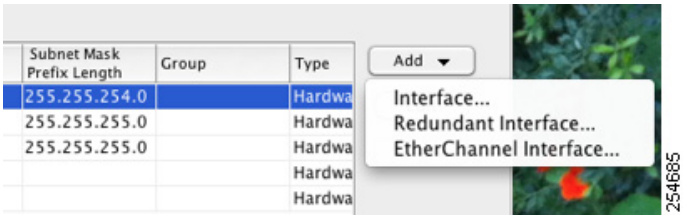
Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Detailed Steps

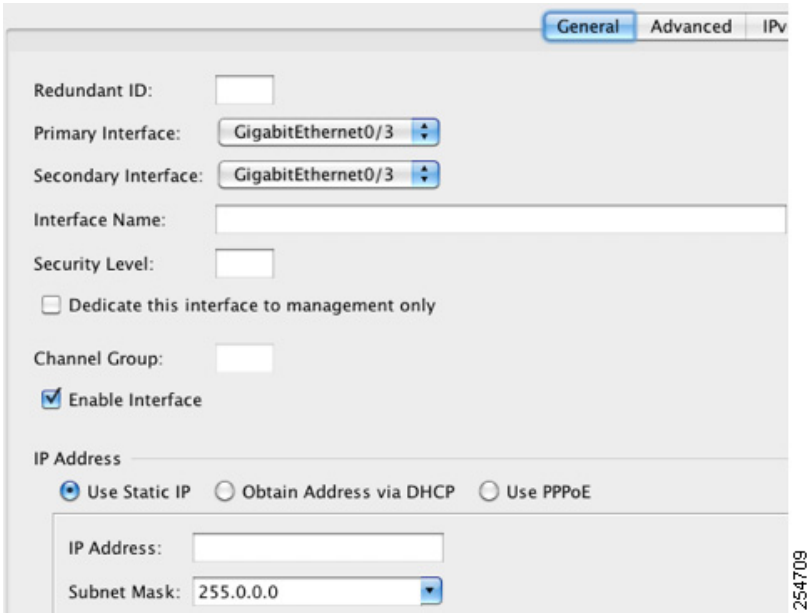
Step 1 Depending on your context mode:

- 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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 8-15.

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 11-36.

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 8-15.
- For single context mode, complete the interface configuration. See the Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 redundant*number* **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 redundant*number* **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 11-31
- Customizing the EtherChannel, page 11-33

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 10-43 instead of this procedure.
- See also the "EtherChannel Guidelines" section on page 11-13.

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.

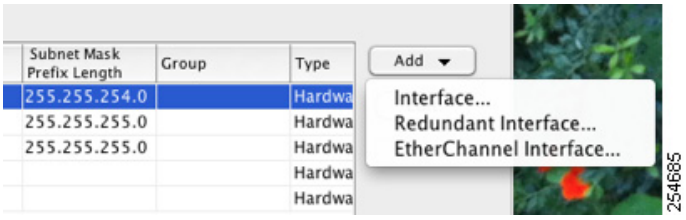
Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Detailed Steps

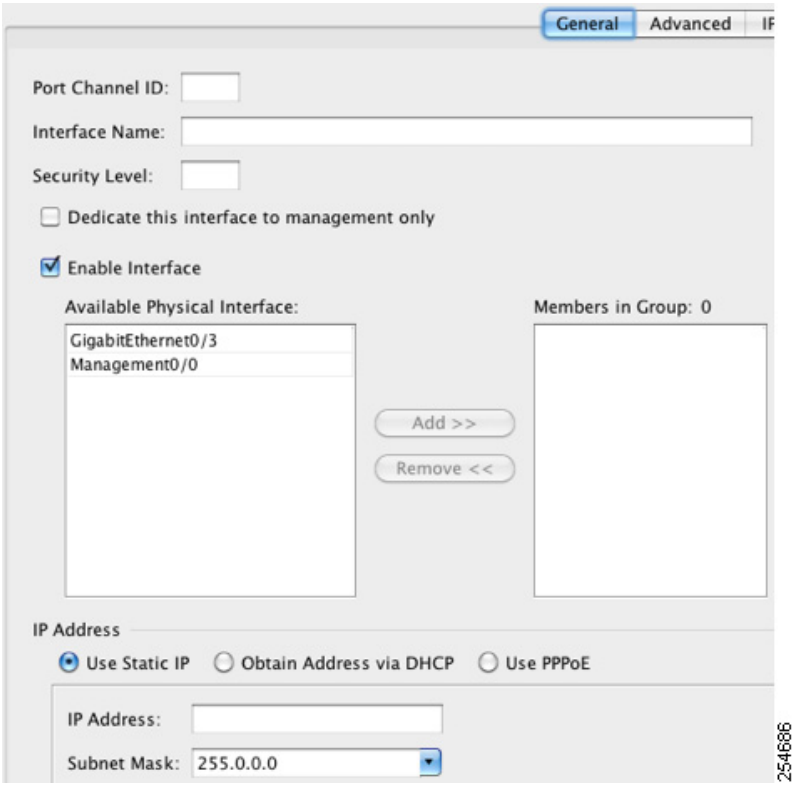
Step 1 Depending on your context mode:

- 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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 8-15.

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.

Note If you want to set the EtherChannel mode to On, then you must include only one interface initially. After you complete this procedure, edit the member interface, and set the mode to On. Apply your changes, then edit the EtherChannel to add more member interfaces.

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.



Step 7 Click **Apply**. All member interfaces are enabled automatically.

What to Do Next

Optional Tasks:

- Customize the EtherChannel interface. See the "Customizing the EtherChannel" section on page 11-33.
- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 11-36.

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 8-15.
- For single context mode, complete the interface configuration. See the Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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

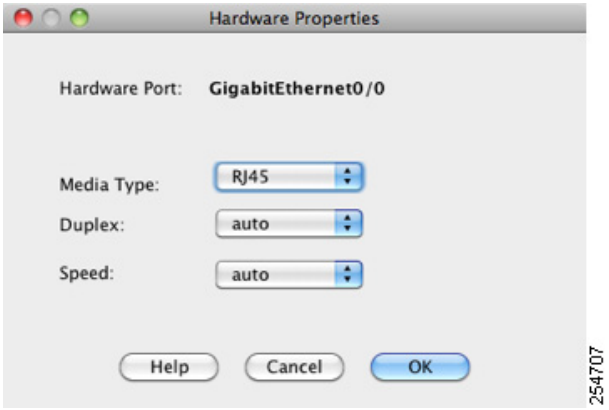
Step 1 Depending on your context mode:

- 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 11-2.

d. To enable pause (XOFF) frames for flow control on 1-Gigabit and 10-Gigabit Ethernet interfaces, check the **Enable Pause Frame** check box.

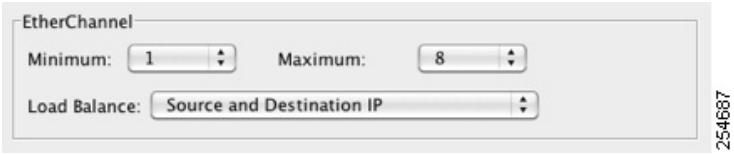
   If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. Pause (XOFF) and XON frames are generated automatically by the NIC hardware based 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 11-7.

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 11-6.

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 11-36.

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 8-15.
- For single context mode, complete the interface configuration. See the Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 11-10.
- 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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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

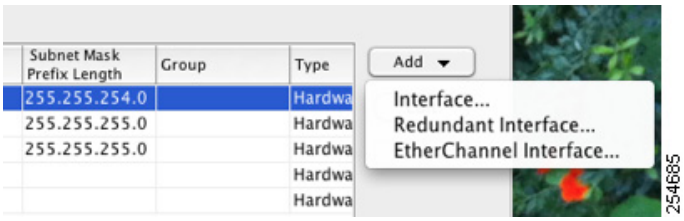
For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Detailed Steps

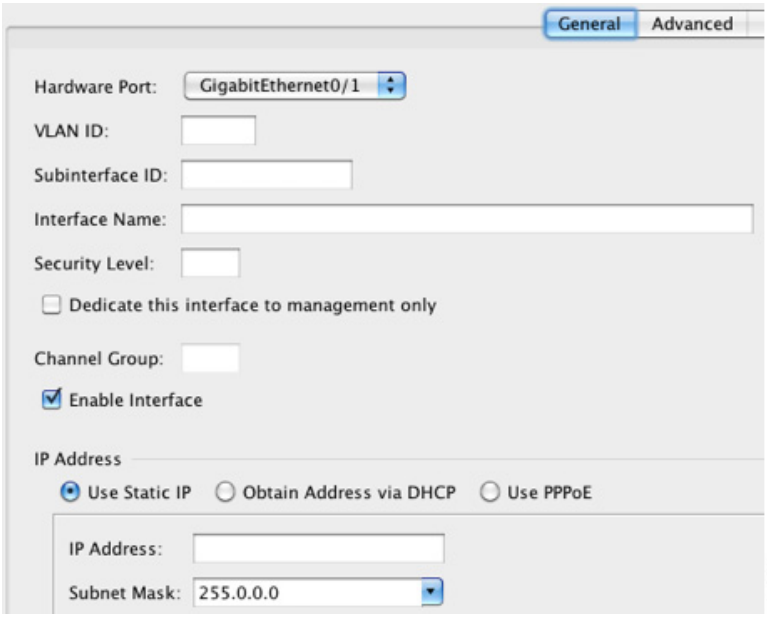
Step 1 Depending on your context mode:

- 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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 8-15.

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 11-39.

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 ACLs. See the "Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size" section on page 11-8 for more information.

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, MTU, and TCP MSS" section on page 13-13. In multiple context mode, set the MTU within each context.
- Be sure to adjust the TCP MSS, either to disable it for non-VPN traffic, or to increase it in accord with the MTU according to the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13.

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 11-40
- MAC Address Table, page 11-40
- Interface Graphs, page 11-40

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 11-40 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 11-40.
- 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 8, "Configuring Multiple Context Mode."
  b. Complete the interface configuration according to Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."
- For single context mode, complete the interface configuration according to Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Feature History for ASA 5510 and Higher Interfaces

Table 11-3 lists the release history for this feature.

Table 11-3 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).
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 ACLs. 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.

Table 11-3 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
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 12-16.

This chapter includes the following sections:

- Information About ASA 5505 Interfaces, page 12-1
- Licensing Requirements for ASA 5505 Interfaces, page 12-4
- Guidelines and Limitations, page 12-5
- Default Settings, page 12-5
- Starting ASA 5505 Interface Configuration, page 12-6
- Monitoring Interfaces, page 12-12
- Where to Go Next, page 12-15
- Feature History for ASA 5505 Interfaces, page 12-16

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 12-2
- Maximum Active VLAN Interfaces for Your License, page 12-2
- VLAN MAC Addresses, page 12-4
- Power over Ethernet, page 12-4
- Monitoring Traffic Using SPAN, page 12-4
- Auto-MDI/MDIX Feature, page 12-4

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 12-4 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 12-1 for more information.
- Security Plus license—20 active VLANs.

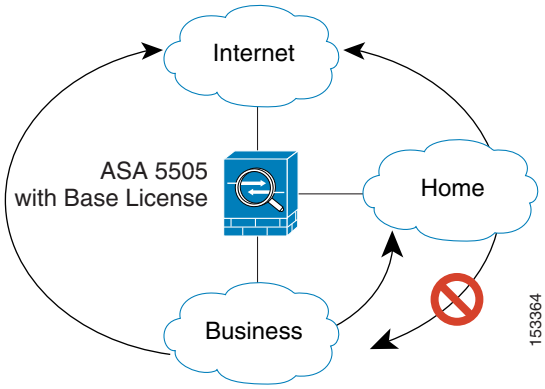
In transparent firewall mode, you can configure the following VLANs depending on your license:

- 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 12-1 for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 12-1 ASA 5505 with Base License

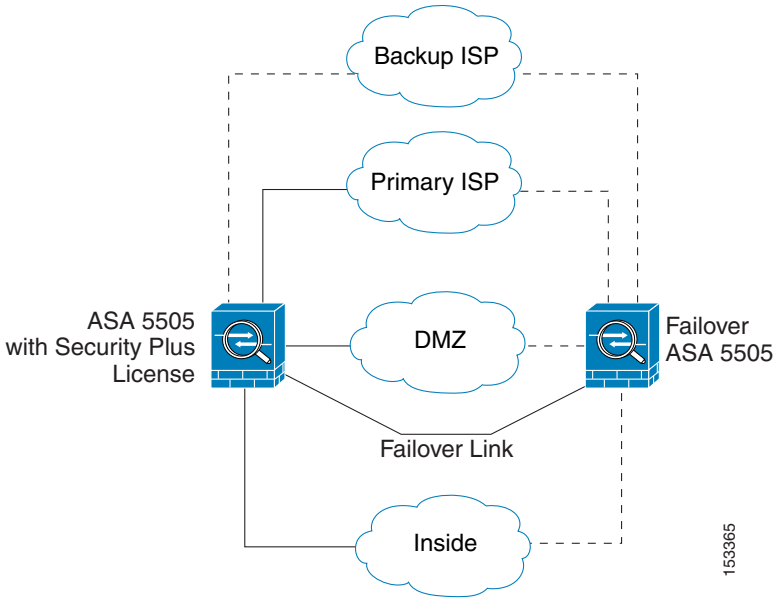


With the Security Plus license, you can configure 20 VLAN interfaces in routed mode, including a VLAN interface for failover and a VLAN interface as a backup link to your ISP. You can configure the backup interface to not pass through traffic unless the route through the primary interface fails. You can configure trunk ports to accommodate multiple VLANs per port.

Note The ASA 5505 supports Active/Standby failover, but not Stateful Failover.

See Figure 12-2 for an example network.

Figure 12-2 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, MTU, and TCP MSS" section on page 13-13.
- 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, MTU, and TCP MSS" section on page 14-15.

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 12-8 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: Routed Mode: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 Transparent Mode: 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. VLAN Trunks: Base License: None. Security Plus License: 8.

Guidelines and Limitations

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.

Failover Guidelines

Active/Standby failover is only supported with the Security Plus license. Active/Active failover is not supported.

IPv6 Guidelines

Supports IPv6.

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 3-19.

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 12-6
- Configuring VLAN Interfaces, page 12-6
- Configuring and Enabling Switch Ports as Access Ports, page 12-8
- Configuring and Enabling Switch Ports as Trunk Ports, page 12-10

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 12-6.

Step 2 Configure and enable switch ports as access ports. See the "Configuring and Enabling Switch Ports as Access Ports" section on page 12-8.

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 12-10.

Step 4 Complete the interface configuration according to Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 12-1.

Guidelines

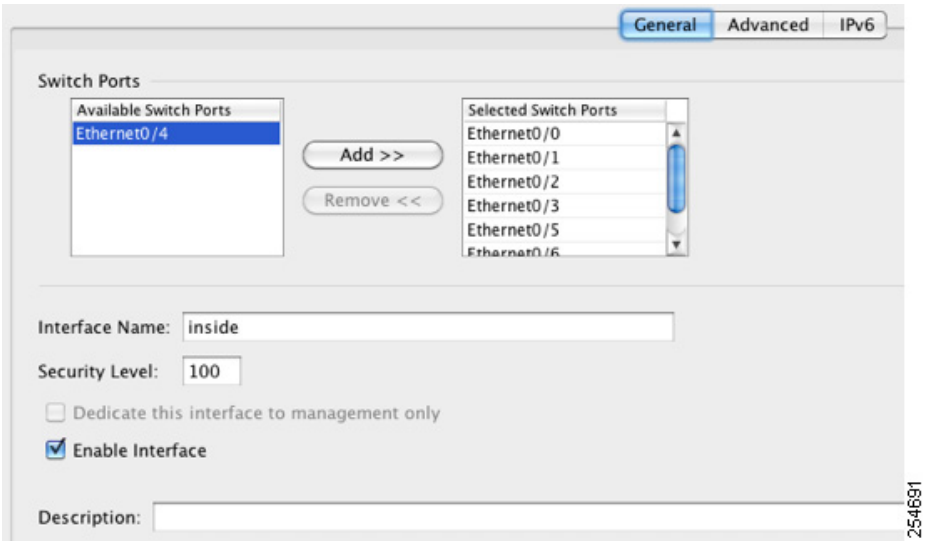
We suggest that you finalize your interface configuration before you enable Easy VPN. If you enabled Easy VPN, you cannot add or delete VLAN interfaces, nor can you edit the security level or interface name.

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 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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, MTU, and TCP MSS" section on page 13-13.

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 12-8 and the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 12-10.

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 12-10. If you have a factory default configuration, see the "ASA 5505 Default Configuration" section on page 3-21 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 12-1.

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.

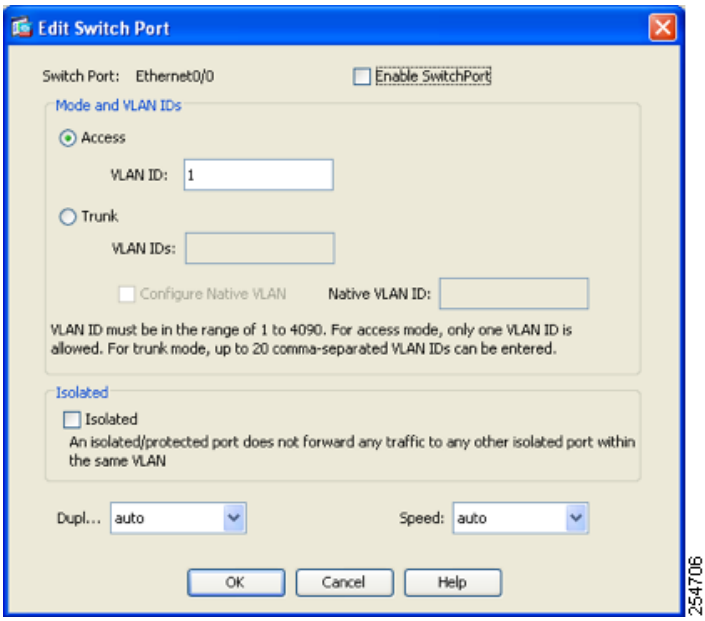
Detailed Steps

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 12-6 (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 12-6 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 12-6 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 12-10.
- To complete the interface configuration, see Chapter 13, "Completing Interface Configuration (Routed Mode)," or Chapter 14, "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 12-8.

Guidelines

This switch port cannot pass traffic until you assign at least one VLAN to it, native or non-native.

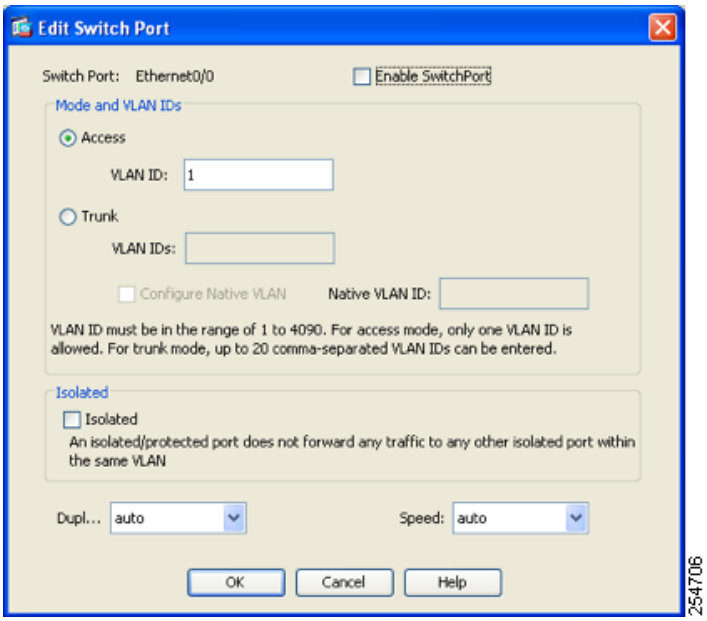
Detailed Steps

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 12-6 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 12-6 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 12-12
- MAC Address Table, page 12-12
- Interface Graphs, page 12-13

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 12-12 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 12-12.
- 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 13, “Completing Interface Configuration (Routed Mode),” or Chapter 14, “Completing Interface Configuration (Transparent Mode, 8.4 and Later).”

Feature History for ASA 5505 Interfaces

Table 12-1 lists the release history for this feature.

Table 12-1 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 13-1
- Licensing Requirements for Completing Interface Configuration in Routed Mode, page 13-2
- Guidelines and Limitations, page 13-5
- Default Settings, page 13-6
- Completing Interface Configuration in Routed Mode, page 13-6
- Turning Off and Turning On Interfaces, page 13-22
- Monitoring Interfaces, page 13-22
- Feature History for Interfaces in Routed Mode, page 13-30

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 13-1
- Dual IP Stack (IPv4 and IPv6), page 13-2

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 13-20 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 ACL to the interface.

  If you enable communication for same security interfaces (see the "Allowing Same Security Level Communication" section on page 13-20), 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: Routed Mode: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 Transparent Mode: 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. VLAN Trunks: Base License: None. Security Plus License: 8.

Model	License Requirement
ASA 5510	VLANs[1]: 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[2]: Base License: 364 Security Plus License: 564
ASA 5520	VLANs[1]: Base License: 150. Interfaces of all types[2]: Base License: 764
ASA 5540	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 964
ASA 5550	VLANs[1]: Base License: 400 Interfaces of all types[2]: Base License: 1764

Model	License Requirement
ASA 5580	VLANs[1]: Base License: 1024 Interfaces of all types[2]: Base License: 4612
ASA 5512-X	VLANs[1]: Base License: 50 Security Plus License: 100 Interfaces of all types[2]: Base License: 716 Security Plus License: 916
ASA 5515-X	VLANs[1]: Base License: 100 Interfaces of all types[2]: Base License: 916
ASA 5525-X	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 1316
ASA 5545-X	VLANs[1]: Base License: 300 Interfaces of all types[2]: Base License: 1716
ASA 5555-X	VLANs[1]: Base License: 500 Interfaces of all types[2]: Base License: 2516
ASA 5585-X	VLANs[1]: Base and Security Plus 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[2]: Base and Security Plus License: 4612

1. For an interface to count against the VLAN limit, you must assign a VLAN to it.
2. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** defined in the configuration counts against this limit.

Model	License Requirement
ASA SM	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 11, "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 8-15.
- PPPoE is not supported in multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode. For transparent mode, see Chapter 14, "Completing Interface Configuration (Transparent Mode, 8.4 and Later)."

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See Chapter 9, "Configuring Failover," 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 3-19.

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, MTU, and TCP MSS" section on page 13-13.

Completing Interface Configuration in Routed Mode

This section includes the following topics:

- Task Flow for Completing Interface Configuration, page 13-7
- Configuring General Interface Parameters, page 13-7
- Configuring the MAC Address, MTU, and TCP MSS, page 13-13
- Configuring IPv6 Addressing, page 13-15
- Allowing Same Security Level Communication, page 13-20

Task Flow for Completing Interface Configuration

Step 1 Set up your interfaces depending on your model:

- ASA 5510 and higher—Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)."
- ASA 5505—Chapter 12, "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 8-15.

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 13-7.

Step 5 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13.

Step 6 (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 13-15.

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 13-20.

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 Chapter 9, "Configuring Failover," 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 11, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 12, "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 8-15.
- 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.

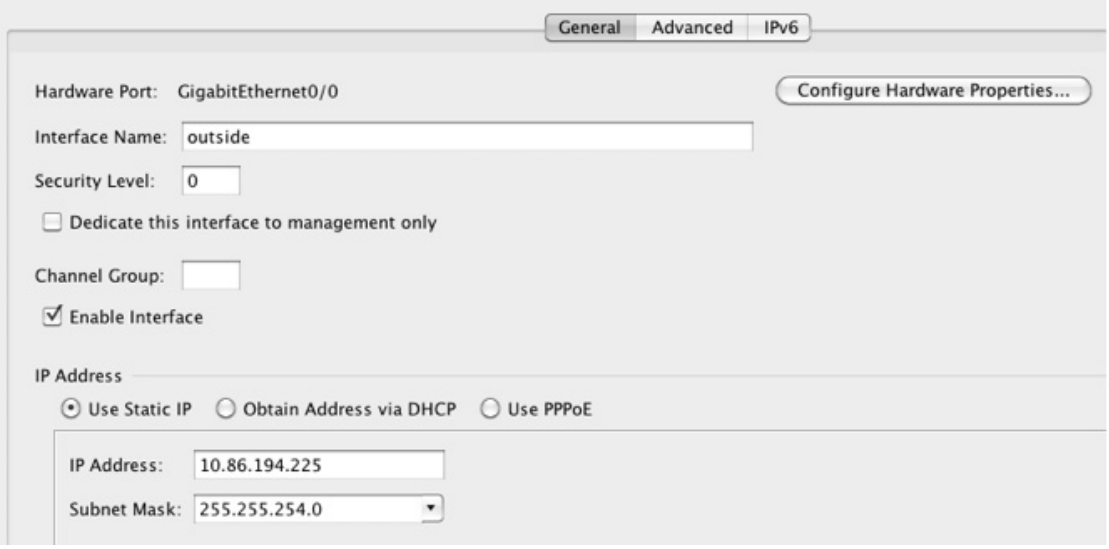
Detailed Steps

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 13-1 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 13-8 for more information.

(ASA 5512-X through ASA 5555-X) You cannot disable this option on the Management 0/0 interface.

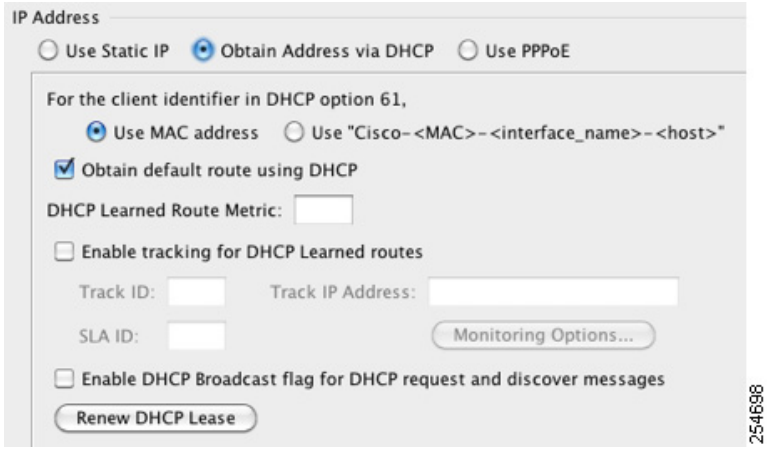
Note The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

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.

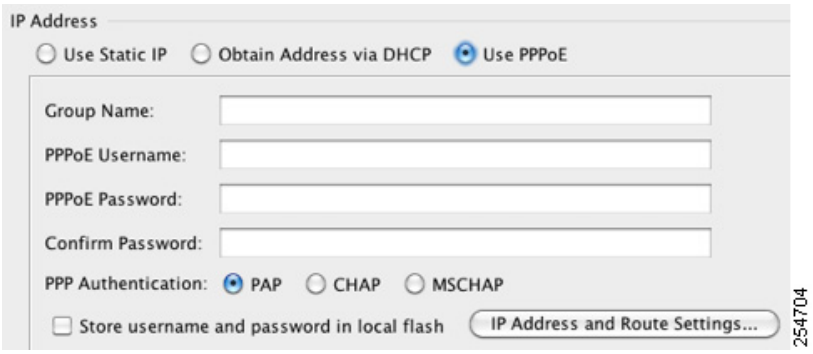
Monitor Options—Click this button to open the Route Monitoring Options dialog box. In the Route Monitoring Options dialog box you can configure the parameters of the tracked object monitoring process.

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 13-11 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 11-25.

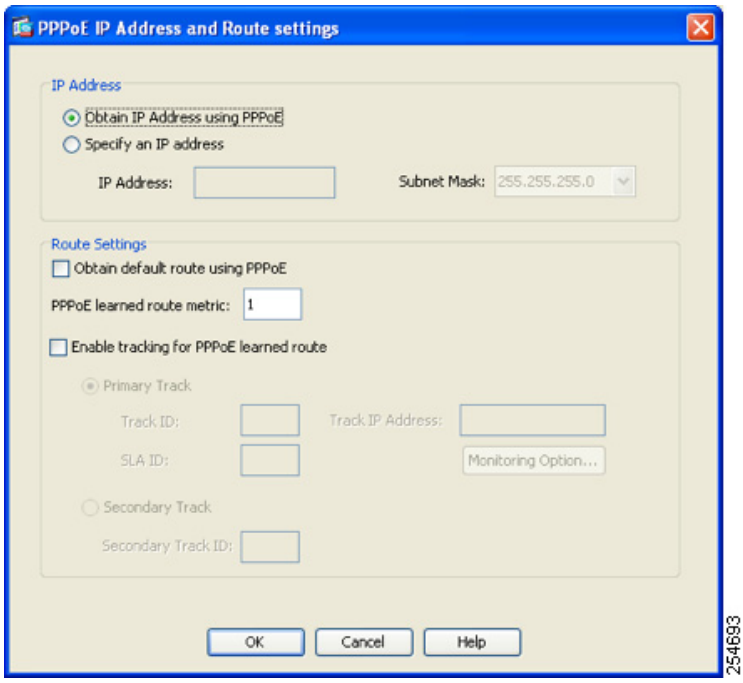
Step 9 Click **OK**.

What to Do Next

- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 13-13.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 13-15.

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, MTU, and TCP MSS

This section describes how to configure MAC addresses for interfaces, how to set the MTU, and set the TCP MSS.

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 8-3 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 8-24 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 and TCP MSS

See the "Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size" section on page 11-8.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 12, "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 8-15.
- 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

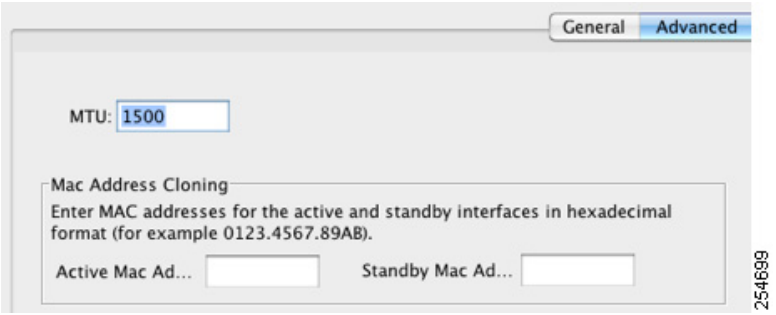
Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

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 11-39.

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.

Step 7 To set the TCP MSS, choose **Configuration > Firewall > Advanced > TCP Options**. Set the following options:

- 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.
- 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).

What to Do Next

(Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 13-15.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the "IPv6 Addresses" section on page 48-5.

This section includes the following topics:

- Information About IPv6, page 13-15
- Configuring a Global IPv6 Address, page 13-16
- Configuring IPv6 Neighbor Discovery, page 13-18
- (Optional) Configuring the Link-Local Addresses Automatically, page 13-18
- (Optional) Configuring the Link-Local Addresses Manually, page 13-19

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 13-16
- Modified EUI-64 Interface IDs, page 13-16

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 11, “Starting Interface Configuration (ASA 5510 and Higher).”

  - ASA 5505—Chapter 12, "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 8-15.
- 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.

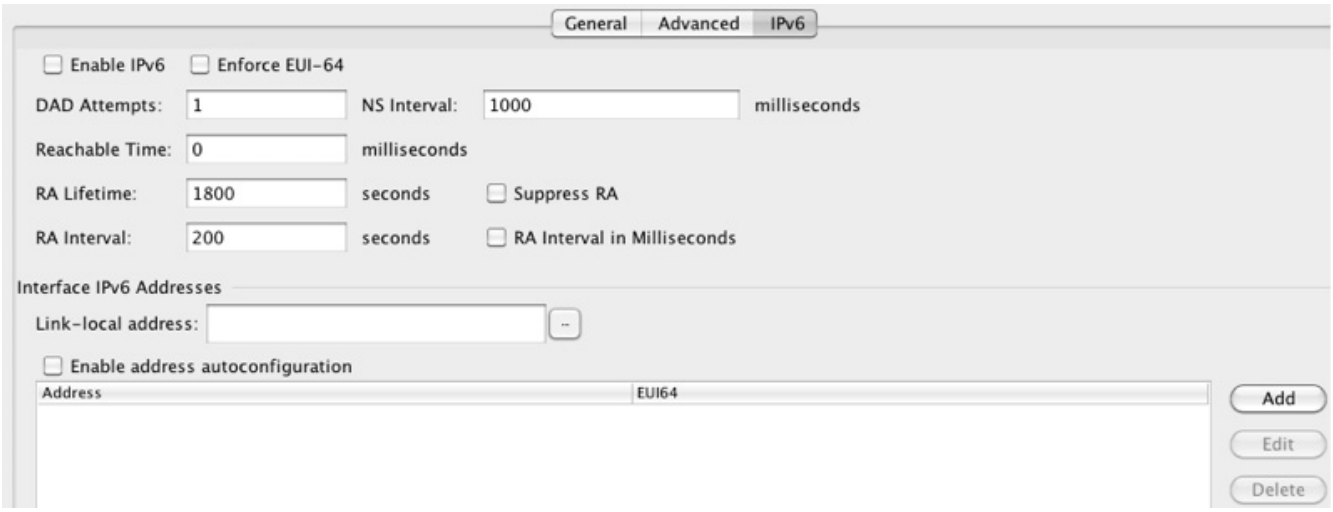
Detailed Steps

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 13-16 for more information.

Step 6 (Optional) In the top area, customize the IPv6 configuration by referring to Chapter 31, "Configuring IPv6 Neighbor Discovery."

Step 7 Configure the global IPv6 address using one of the following methods.

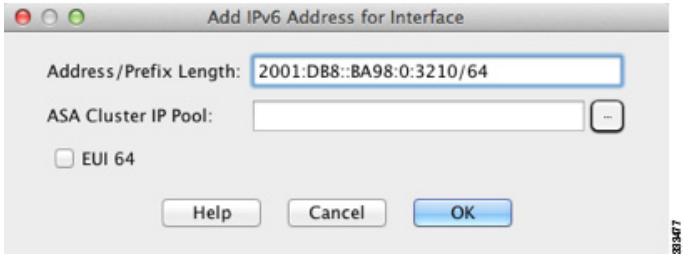
- Stateless autoconfiguration—In the Interface IPv6 Addresses area, check the **Enable address autoconfiguration** check box.

  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 48-5 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 10-38.

  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 31-15.

Step 9 Click **OK**.

You return to the Configuration > Device Setup > Interfaces pane.

Configuring IPv6 Neighbor Discovery

See Chapter 31, "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 13-19.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 13-16.

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 13-18.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 13-16.

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 48-5 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 ACLs.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

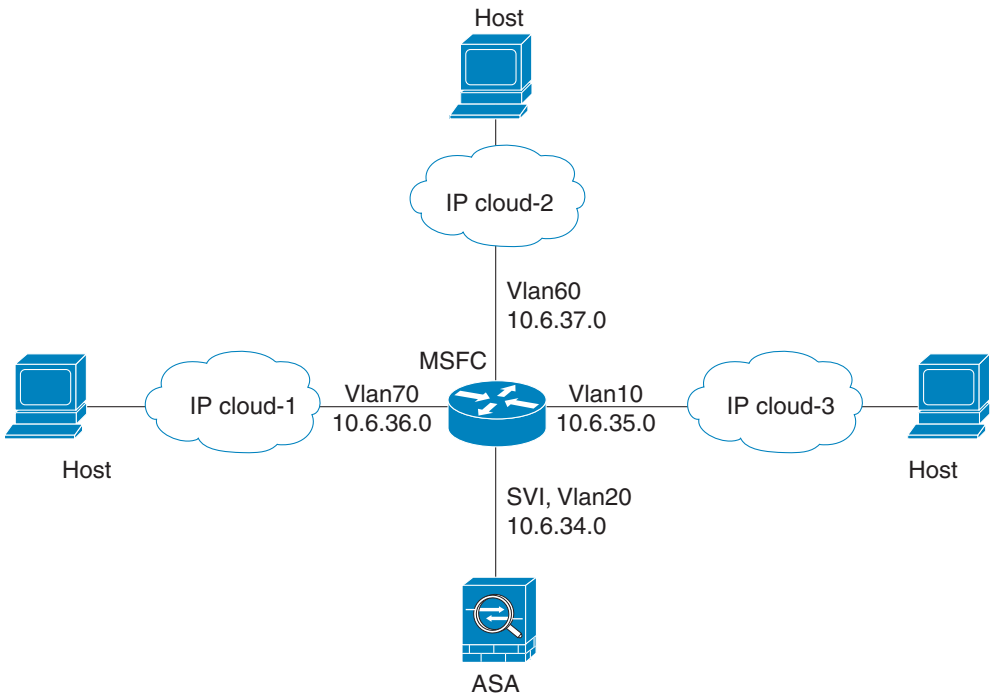
Information About Intra-Interface Communication

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 13-1 shows a network where hosts on the same interface need to communicate.

Figure 13-1 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 13-1:

```
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

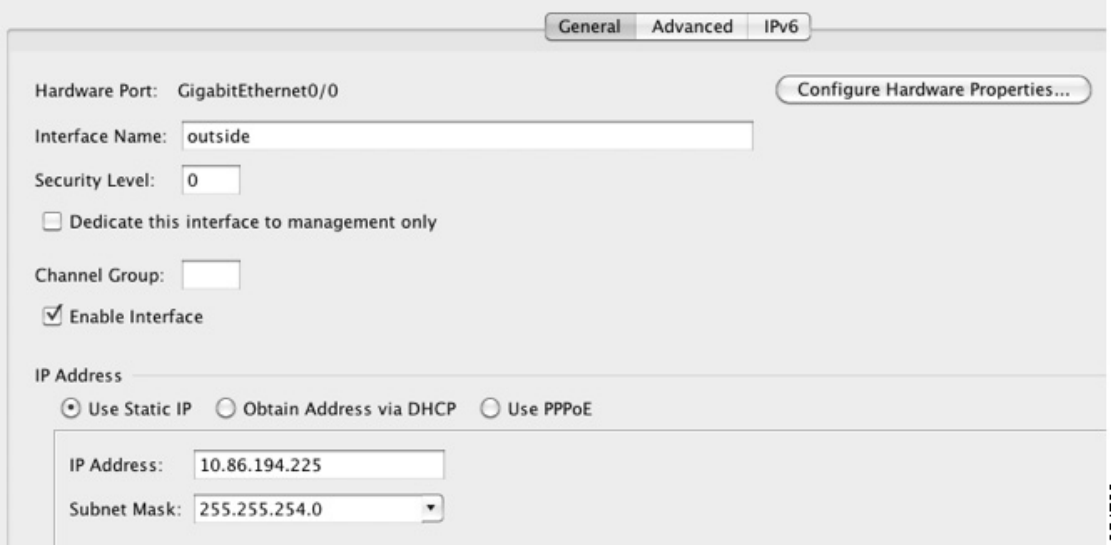
Step 1 Depending on your context mode:

- 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 13-23
- DHCP, page 13-23
- MAC Address Table, page 13-26
- Dynamic ACLs, page 13-26

- Interface Graphs, page 13-26
- PPPoE Client, page 13-29
- Interface Connection, page 13-29

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 13-26 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 13-26.
- 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 13-29
- Monitoring Statistics for, page 13-29

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 13-1 lists the release history for this feature.

Table 13-1 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 ACLs. We modified the following screen: Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced.

Table 13-1 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 15, "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 14-1
- Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 14-3
- Guidelines and Limitations, page 14-5
- Default Settings, page 14-7
- Completing Interface Configuration in Transparent Mode (8.4 and Later), page 14-7
- Turning Off and Turning On Interfaces, page 14-22
- Monitoring Interfaces, page 14-22
- Feature History for Interfaces in Transparent Mode, page 14-30

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 14-2
- Security Levels, page 14-2

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 14-21 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 ACL to the interface.

  If you enable communication for same security interfaces (see the "Allowing Same Security Level Communication" section on page 14-21), 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: Routed Mode: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 Transparent Mode: 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. VLAN Trunks: Base License: None. Security Plus License: 8.

Model	License Requirement
ASA 5510	VLANs[1]: 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[2]: Base License: 364 Security Plus License: 564
ASA 5520	VLANs[1]: Base License: 150. Interfaces of all types[2]: Base License: 764
ASA 5540	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 964

Model	License Requirement
ASA 5550	VLANs[1]: Base License: 400 Interfaces of all types[2]: Base License: 1764
ASA 5580	VLANs[1]: Base License: 1024 Interfaces of all types[2]: Base License: 4612
ASA 5512-X	VLANs[1]: Base License: 50 Security Plus License: 100 Interfaces of all types[2]: Base License: 716 Security Plus License: 916
ASA 5515-X	VLANs[1]: Base License: 100 Interfaces of all types[2]: Base License: 916
ASA 5525-X	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 1316
ASA 5545-X	VLANs[1]: Base License: 300 Interfaces of all types[2]: Base License: 1716

Model	License Requirement
ASA 5555-X	VLANs[1]: Base License: 500 Interfaces of all types[2]: Base License: 2516
ASA 5585-X	VLANs[1]: Base and Security Plus 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[2]: Base and Security Plus License: 4612

1. For an interface to count against the VLAN limit, you must assign a VLAN to it.
2. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** defined in the configuration counts against this limit.

Model	License Requirement
ASA SM	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 11, "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 11-2 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 14-8 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 Chapter 9, "Configuring Failover," 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 3-19.

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, MTU, and TCP MSS" section on page 14-15.

Completing Interface Configuration in Transparent Mode (8.4 and Later)

This section includes the following topics:

- Task Flow for Completing Interface Configuration, page 14-8
- Configuring Bridge Groups, page 14-8
- Configuring General Interface Parameters, page 14-9
- Configuring a Management Interface (ASA 5510 and Higher), page 14-12
- Configuring the MAC Address, MTU, and TCP MSS, page 14-15
- Configuring IPv6 Addressing, page 14-17
- Allowing Same Security Level Communication, page 14-21

Task Flow for Completing Interface Configuration

Step 1 Set up your interfaces depending on your model:

- ASA 5510 and higher—Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)."
- ASA 5505—Chapter 12, "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 8-15.

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 14-8.

Step 5 Configure general interface parameters, including the bridge group it belongs to, the interface name, and security level. See the "Configuring General Interface Parameters" section on page 14-9.

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 14-12.

Step 7 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 14-15.

Step 8 (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 14-17.

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 14-21.

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

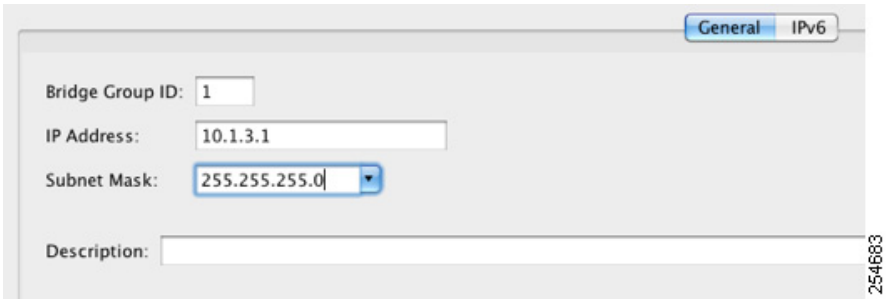
You can configure up to 8 bridge groups in single mode or per context in multiple mode. Note that you must use at least one bridge group; data interfaces must belong to a bridge group.

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.

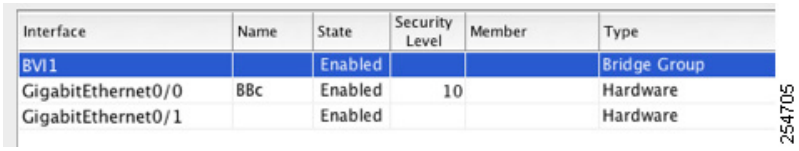
Step 4 In the Subnet Mask field, enter the subnet mask or choose one from the menu.

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 14-9.

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 14-12.

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 14-2.
- If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See Chapter 9, "Configuring Failover," to configure the failover and state links.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 12, "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 8-15.
- 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.

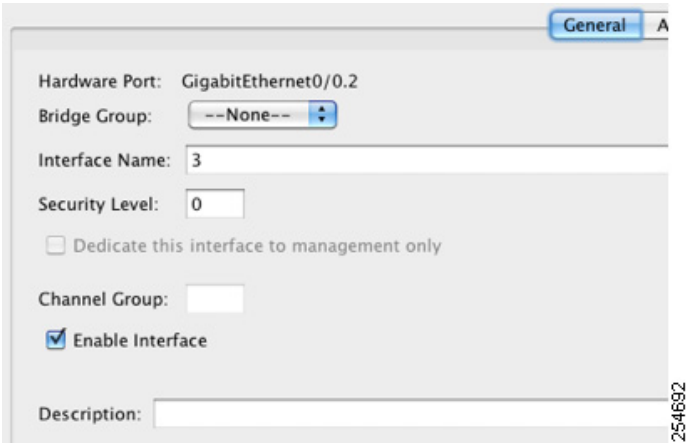
Detailed Steps

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 14-12 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 14-2 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 14-12 for this option.

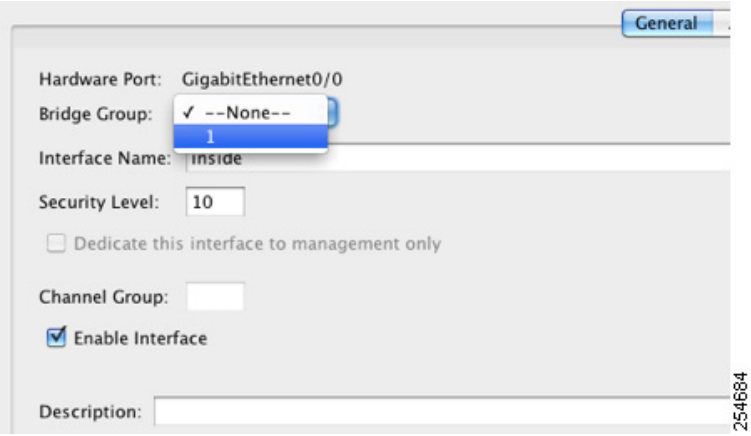
Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

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 11-25.



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 14-12.
- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 14-15.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 14-17.

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 11-2.

Restrictions

- See the "Management Interface" section on page 11-2.
- 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 11, "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 8-15.
- 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.

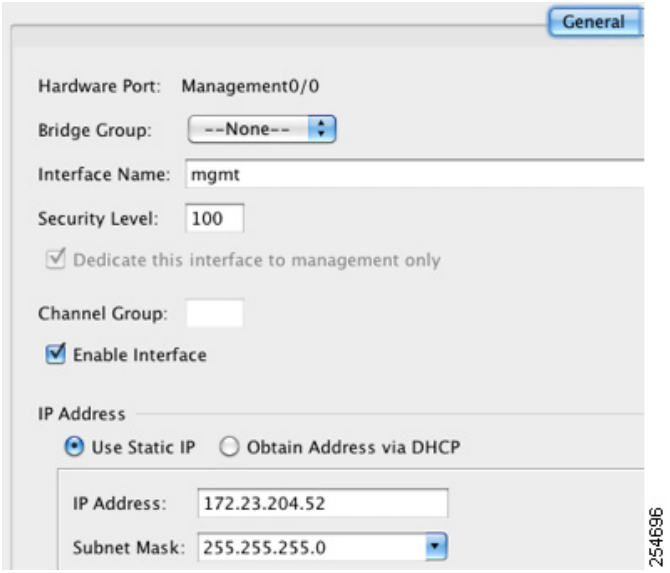
Detailed Steps

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 14-2 for more information.

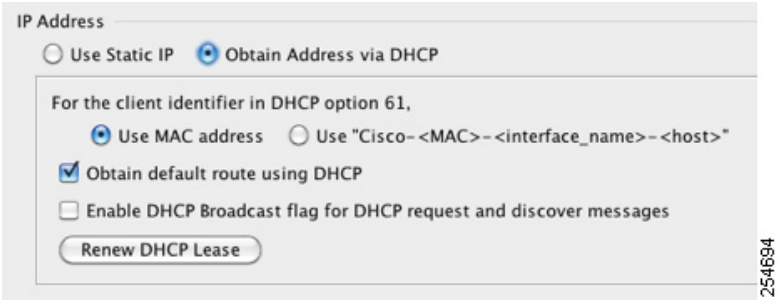
Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

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 11-25.

Step 9 Click **OK**.

What to Do Next

- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 14-15.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 14-17.

Configuring the MAC Address, MTU, and TCP MSS

This section describes how to configure MAC addresses for interfaces, how to set the MTU, and set the TCP MSS.

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 8-3 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 8-24 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 and TCP MSS

See the "Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size" section on page 11-8.

Prerequisites

- Set up your interfaces depending on your model:
  - ASA 5510 and higher—Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 12, "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 8-15.
- 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

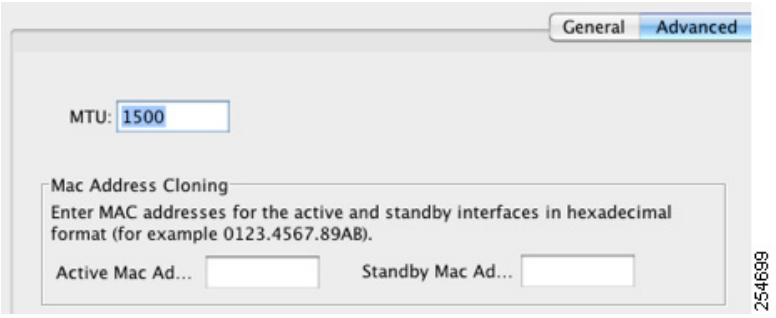
Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

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 11-39.

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.

Step 7 To set the TCP MSS, choose **Configuration > Firewall > Advanced > TCP Options**. Set the following options:

- 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.
- 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).

What to Do Next

(Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 14-17.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the "IPv6 Addresses" section on page 48-5.

This section includes the following topics:

- Information About IPv6, page 14-17
- Configuring a Global IPv6 Address, page 14-18
- Configuring IPv6 Neighbor Discovery, page 14-20
- (Optional) Configuring the Link-Local Addresses Automatically, page 14-20
- (Optional) Configuring the Link-Local Addresses Manually, page 14-21

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 14-18
- Modified EUI-64 Interface IDs, page 14-18
- Unsupported Commands, page 14-18

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 11, "Starting Interface Configuration (ASA 5510 and Higher)."
  - ASA 5505—Chapter 12, "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 8-15.
- 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.

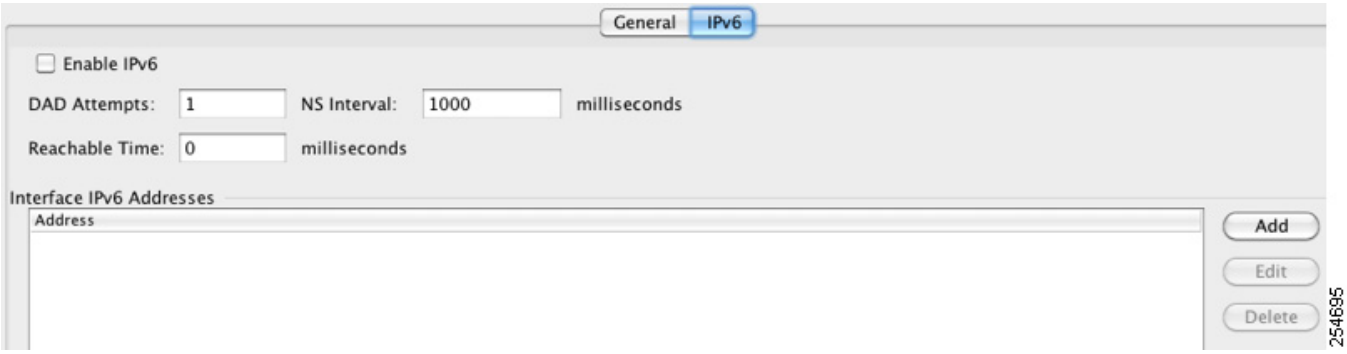
Detailed Steps

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.

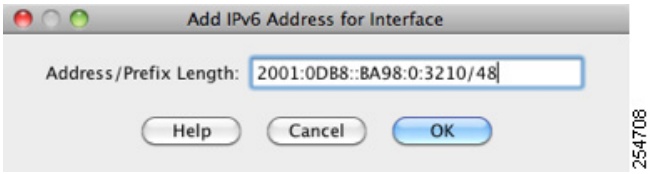
See the "Modified EUI-64 Interface IDs" section on page 14-18 for more information.

Step 6 (Optional) In the top area, customize the IPv6 configuration by referring to Chapter 31, "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 48-5 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 31, "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 14-21.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 14-18.

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 14-20.

For other IPv6 options, including enforcing the Modified EUI-64 format, and DAD settings, see the "Configuring a Global IPv6 Address" section on page 14-18.

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 48-5 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 ACLs.

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

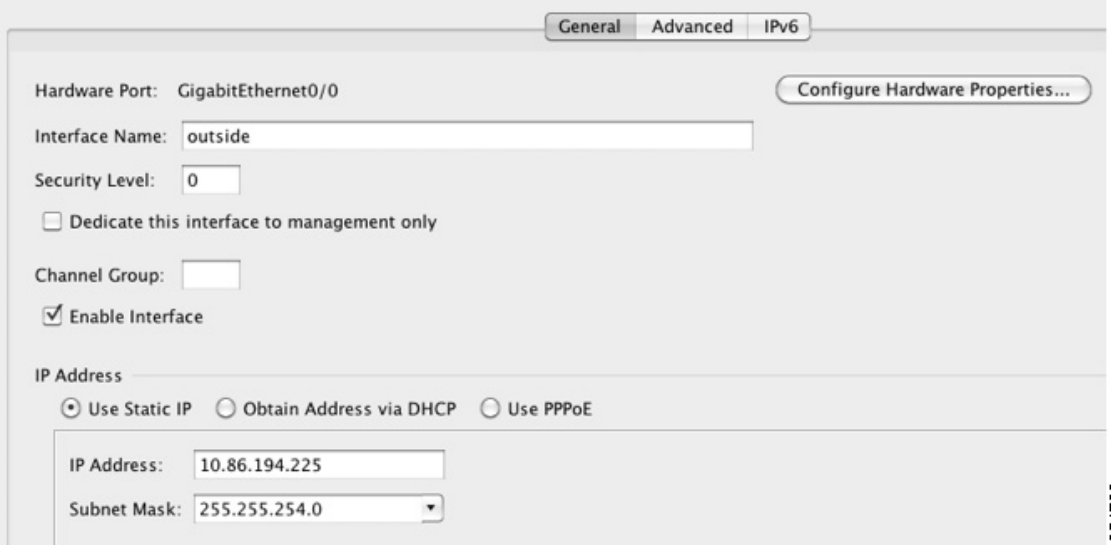
Step 1 Depending on your context mode:

- 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 14-23
- DHCP, page 14-23
- MAC Address Table, page 14-26
- Dynamic ACLs, page 14-26

- Interface Graphs, page 14-26
- PPPoE Client, page 14-29
- Interface Connection, page 14-29

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 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 14-26 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 14-26.
- 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 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 14-29
- Monitoring Statistics for, page 14-29

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 14-1 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 14-1 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 14-1 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 ACLs. 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 15, "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 15-1
- Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 15-3
- Guidelines and Limitations, page 15-3
- Default Settings, page 15-4
- Setting the Management IP Address for a Transparent Firewall (8.3 and Earlier), page 15-4
- Completing Interface Configuration in Transparent Mode (8.3 and Earlier), page 15-9
- Monitoring Interfaces, page 15-17
- Feature History for Interfaces in Transparent Mode, page 15-18

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 15-2
- Security Levels, page 15-2

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 15-9.

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 15-17 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 ACL to the interface.

  If you enable communication for same security interfaces (see the "Allowing Same Security Level Communication" section on page 15-17), 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 11, "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 11-2 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 15-4 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 Chapter 9, "Configuring Failover," 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 3-19.

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 15-4
- Configuring the IPv6 Address, page 15-5

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 15-5
- Configuring the Global Address, page 15-7
- Configuring the Link-Local Addresses Automatically, page 15-7
- Configuring the Link-Local Address on an Interface Manually, page 15-8
- Configuring DAD Settings, page 15-8

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- IPv6 Addressing, page 15-5
- Duplicate Address Detection, page 15-5
- Modified EUI-64 Interface IDs, page 15-6
- Unsupported Commands, page 15-6

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 48-5 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 48-5 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 15-8.

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 15-6 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 48-5 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 15-7.

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 15-9
- Configuring General Interface Parameters, page 15-9
- Configuring a Management Interface (ASA 5510 and Higher), page 15-11
- Configuring the MAC Address, MTU, and TCP MSS, page 15-14
- Allowing Same Security Level Communication, page 15-17

Task Flow for Completing Interface Configuration

Step 1 Complete the procedures in the "Starting Interface Configuration (ASA 5510 and Higher)" section on page 11-14 or the "Starting ASA 5505 Interface Configuration" section on page 12-6.

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 15-9.

Step 4 (Optional) Configure a management interface. See the "Configuring a Management Interface (ASA 5510 and Higher)" section on page 15-11.

Step 5 (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 15-14.

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 15-17.

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 15-11.

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 15-2.
- If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See Chapter 9, "Configuring Failover," to configure the failover and state links.

Prerequisites

- Complete the procedures in Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 12, "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 15-2 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 11-25.

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 15-11.
- (Optional) Configure the MAC address and the MTU. See the "Configuring the MAC Address, MTU, and TCP MSS" section on page 15-14.

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 11-2.

This section includes the following topics:

- Configuring General Parameters and the IPv4 Address, page 15-11
- Configuring a Global IPv6 Address and Other Options, page 15-13

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 11, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 12, "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 15-2 for more information.

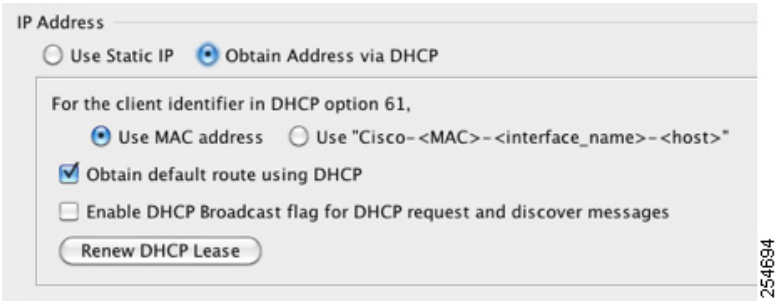
Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

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 11-25.

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 11, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 12, "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.

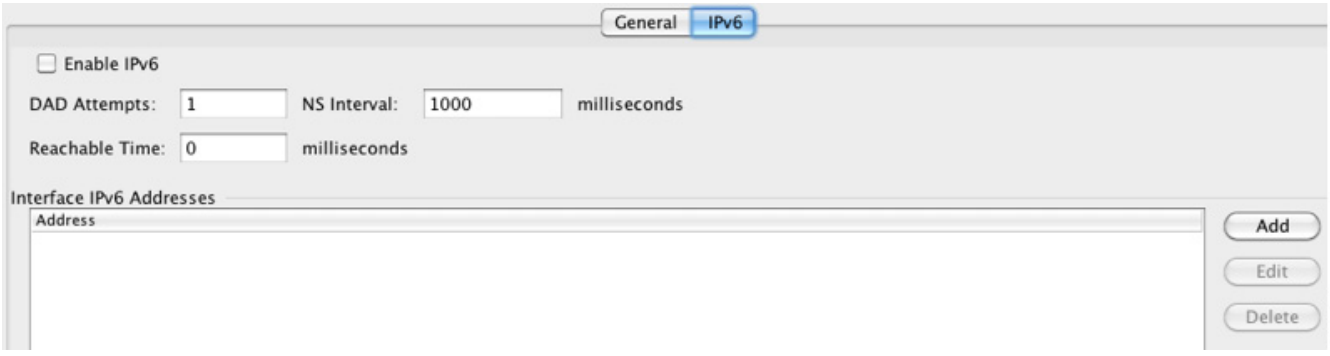
Detailed Steps

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.



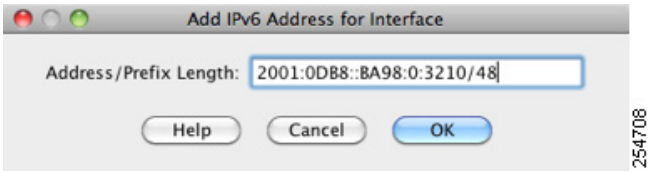
Step 4 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

See the "Modified EUI-64 Interface IDs" section on page 15-6 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 48-5 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, MTU, and TCP MSS" section on page 15-14.

Configuring the MAC Address, MTU, and TCP MSS

This section describes how to configure MAC addresses for interfaces, how to set the MTU, and set the TCP MSS.

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 8-3 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 8-24 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 and TCP MSS

See the "Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size" section on page 11-8.

Prerequisites

- Complete the procedures in Chapter 11, "Starting Interface Configuration (ASA 5510 and Higher)," or Chapter 12, "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

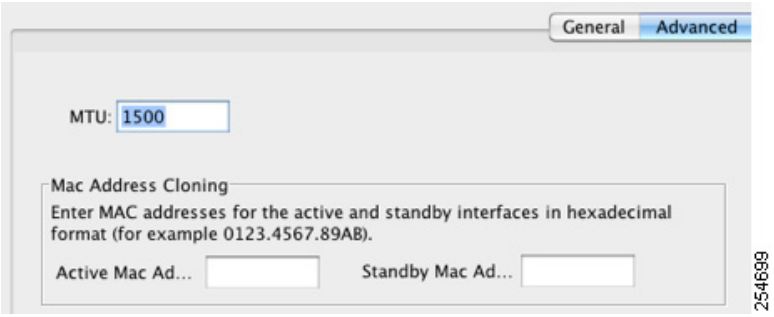
Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

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 (supported models), 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 11-39.

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 ACLs.

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.

Step 7 To set the TCP MSS, choose **Configuration > Firewall > Advanced > TCP Options**. Set the following options:

- 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.

- 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).

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 ACLs.

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 14-22.

Feature History for Interfaces in Transparent Mode

Table 15-1 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 15-1 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 15-1 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 ACLs. 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 14, "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 4

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 16-1
- Setting the Date and Time, page 16-3
- Configuring the Master Passphrase, page 16-5
- Configuring the DNS Server, page 16-8
- Changing the Heap Memory Size, page 16-10
- Monitoring DNS Cache, page 16-10

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 16-1
- Feature History for the Hostname, Domain Name, and Passwords, page 16-3

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.

Detailed Steps

Step 1 In ASDM, choose **Configuration > Device Setup > Device Name/Password**.

Step 2 Enter the hostname. The default hostname is “asa.”

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 with a blank username. (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 for Telnet access.

The Telnet password sets the login password. 9.1(1): By default, it is “cisco.” 9.1(2) and later: There is no default password. The login password lets you access EXEC mode if you connect to the ASA using a Telnet session. (If you configure user authentication for Telnet 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.

Feature History for the Hostname, Domain Name, and Passwords

Table 16-1 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 16-1 Feature History for the Master Passphrase

Feature Name	Platform Releases	Feature Information
Removal of the default Telnet password	9.0(2), 9.1(2)	To improve security for management access to the ASA, the default login password for Telnet was removed; you must manually set the password before you can log in using Telnet. **Note**: The login password is only used for Telnet if you do not configure Telnet user authentication. Formerly, when you cleared the password, the ASA restored the default of "cisco." Now when you clear the password, the password is removed. The login password is also used for Telnet sessions from the switch to the ASASM (see the **session** command). For initial ASASM access, you must use the **service-module session** command, until you set a login password. We did not modify any ASDM screens.

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 16-3
- Setting the Date and Time Manually, page 16-4

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 16-5
- Licensing Requirements for the Master Passphrase, page 16-6
- Guidelines and Limitations, page 16-6
- Adding or Changing the Master Passphrase, page 16-6
- Disabling the Master Passphrase, page 16-7
- Recovering the Master Passphrase, page 16-8
- Feature History for the Master Passphrase, page 16-8

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 16-8.

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 16-2 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 16-2 Feature History for the Master Passphrase

Feature Name	Platform Releases	Feature Information
Master Passphrase	8.3(1)	We introduced this feature. 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. 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 18-3.

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 24-1 for more information about routing.

To configure the DNS server, perform the following steps:

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

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

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

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 DHCP Services

This chapter describes how to configure the DHCP server or DHCP relay and includes the following sections:

- Information About DHCP Services, page 17-1
- Licensing Requirements for DHCP, page 17-2
- Guidelines and Limitations, page 17-2
- Configuring DHCP Services, page 17-4
- Additional References, page 17-9
- Monitoring DHCP Services, page 17-9
- Feature History for DHCP Services, page 17-10

Information About DHCP Services

- Information About the DHCP Server, page 17-1
- Information About the DHCP Relay Agent, page 17-2

Information About the DHCP Server

DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The ASA can provide a DHCP server to DHCP clients attached to ASA interfaces. The DHCP server provides network configuration parameters directly to DHCP clients.

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.

An IPv4 DHCP client uses a broadcast rather than a multicast address to reach the server. The DHCP client listens for messages on UDP port 68; the DHCP server listens for messages on UDP port 67.

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.

Information About the DHCP Relay Agent

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the ASA because it does not forward broadcast traffic.

You can remedy this situation by configuring the interface of your ASA that is receiving the broadcasts to forward DHCP requests to a DHCP server on another interface.

Licensing Requirements for DHCP

Table 17-1 shows the licensing requirements for DHCP.

Table 17-1 Licensing Requirements

Model	License Requirement
All models	Base License.

For all ASA models, 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.

Guidelines and Limitations

Firewall Mode Guidelines

Supported in routed firewall mode.

Not supported in transparent firewall mode. See the "DHCP Relay Guidelines" section on page 17-4 for more information.

Context Mode Guidelines

Supported in single and multiple context mode.

Failover Guidelines

Supports Active/Active and Active/Standby failover.

IPv6 Guidelines

Supports IPv6, except for interface-specific DHCP relay servers.

DHCP Server Guidelines

- The maximum available DHCP pool is 256 addresses.
- 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.
- (9.1.3 and earlier) 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.
- (9.1.4 and later) When a client connects, the ASA sends a discovery message to all the servers in the server list. This message includes the IP address (within a subnetwork) that was configured with the **dhcp-network-scope** command in the group policy. The ASA selects the first offer received and drops the other offers. 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. When the address needs to be renewed, it attempts to renew it with the lease server (the server from which the address was acquired). If the DHCP renew fails after a specified number of retries (four attempts), the ASA moves to the DHCP rebind phase after a predefined time period. During the rebind phase, the ASA simultaneously sends requests to all servers in the group. In a high availability environment, lease information is shared, so the other servers can acknowledge the lease and ASA will return to the bound state. During the rebind phase, if there is no response from any of the servers in the server list (after three retries), then the ASA will purge the entries.

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.

The **dhcp-network-scope** command setting applies only to VPN users.

DHCP Relay Guidelines

- You can configure a maximum of 10 DHCPv4 relay servers in single mode and per context, global and interface-specific servers combined, with a maximum of 4 servers per interface.
- You can configure a maximum of 10 DHCPv6 relay servers in single mode and per context. Interface-specific servers for IPv6 are not supported.
- The relay agent cannot be enabled if the DHCP server feature is also enabled.
- 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, ICMP unreachable, 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.
- DHCP relay services are not available in transparent firewall mode. You can, however, allow DHCP traffic through using an access list. To allow DHCP requests and replies through the ASA in transparent mode, you need to 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.
- For IPv4, clients must be directly-connected to the ASA and cannot send requests through another relay agent or a router. For IPv6, the ASA supports packets from another relay server.
- For multiple context mode, you cannot enable DHCP relay on an interface that is used by more than one context.
- The DHCP clients must be on different interfaces from the DHCP servers to which the ASA relays requests.

Configuring DHCP Services

- Configuring the DHCP Server, page 17-4
- Configuring the DHCP Relay Agent, page 17-7

Configuring the DHCP Server

This section describes how to configure a DHCP server provided by the ASA and includes the following topics:

- Enabling the DHCP Server, page 17-5
- Configuring Advanced DHCP Options, page 17-6

Enabling the DHCP Server

To enable the DHCP server on an ASA interface, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration > Device Management > DHCP > DHCP Server**.

Step 2 Select an interface, and click **Edit**.

a. To enable the DHCP server on the selected interface, check the **Enable DHCP Server** check box.

b. 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.

c. In the Optional Parameters area, set the following:

- The DNS servers (1 and 2) configured for the interface.
- The WINS servers (primary and secondary) configured for the interface.
- The domain name of the interface.
- The time in milliseconds that the ASA will wait for an ICMP ping response on the interface.
- The duration of time that the DHCP server configured on the interface allows DHCP clients to use an assigned IP address.
- The interface on a DHCP client that provides DNS, WINS, and domain name information for automatic configuration if the ASA is acting as a DHCP client on a specified interface (usually outside).
- 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 17-6.

d. In the Dynamic Settings for DHCP Server area, 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

e. Click **OK** to close the Edit DHCP Server dialog box.

Step 3 In the Global DHCP Options area below the DHCP Server table, check the **Enable Auto-configuration from interface** check box to enable DHCP auto configuration only if the ASA is acting as a DHCP client on a specified interface (usually outside).

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 4 Choose the interface from the drop-down list.

Step 5 To override the interface DHCP or PPPoE client WINS parameter with the VPN client parameter, check the **Allow VPN override** check box.

Step 6 In the DNS Server 1 field, enter the IP address of the primary DNS server for a DHCP client.

Step 7 In the DNS Server 2 field, enter the IP address of the alternate DNS server for a DHCP client.

Step 8 In the Domain Name field, enter the DNS domain name for DHCP clients (for example, example.com).

Step 9 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 10 In the Primary WINS Server field, enter the IP address of the primary WINS server for a DHCP client.

Step 11 In the Secondary WINS Server field, enter the IP address of the alternate WINS server for a DHCP client.

Step 12 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 13 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 17-6.

Step 14 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 15 Click **Apply** to save your changes.

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.

Detailed Steps

Step 1 Choose **Configuration > Device Management > DHCP > DHCP Server**, and click **Advanced**.

Step 2 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.

Step 3 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.

Step 4 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 5 Click **OK** to close the Advanced DHCP Options dialog box.

Step 6 Click **Apply** to save your changes.

Configuring the DHCP Relay Agent

When a DHCP request enters an interface, the DHCP servers to which the ASA relays the request depends on your configuration. You can configure the following types of servers:

- Interface-specific DHCP servers—When a DHCP request enters a particular interface, then the ASA relays the request only to the interface-specific servers.
- Global DHCP servers—When a DHCP request enters an interface that does not have interface-specific servers configured, the ASA relays the request to all global servers. If the interface has interface-specific servers, then the global servers are not used.

Detailed Steps

Step 1 Choose **Configuration > Device Management > DHCP > DHCP Relay**.

Step 2 In the DHCP Relay Agent area, check the check boxes for the services you want for each interface:

- **IPv4 > DHCP Relay Enabled**.
- **IPv4 > Set Route**— Changes the default gateway address in the DHCP message from the server to that of the ASA interface that is closest to the DHCP client, which relayed the original DHCP request. This action allows the client to set its default route to point to the ASA even if the DHCP server specifies a different router. If there is no default router option in the packet, the ASA adds one containing the interface address.
- **IPv6 > DHCP Relay Enabled**.
- **Trusted Interface**—Specifies a DHCP client interface that you want to trust. You can configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. You can alternatively trust all interfaces using the **Set dhcp relay information as trusted on all interfaces** check box (see Step 7).

Step 3 In the Global DHCP Relay Servers area, add one or more DHCP servers to which DHCP requests are relayed:

a. Click **Add**. The Add Global DHCP Relay Server dialog box appears.

b. In the DHCP Server field, enter the IPv4 or IPv6 address of the DHCP server.

c. From the Interface drop-down list, choose the interface to which the specified DHCP server is attached.

d. Click **OK**.

The newly added global DHCP relay server appears in the Global DHCP Relay Servers list.

Step 4 (Optional) In the IPv4 Timeout field, enter the amount of time, in seconds, allowed for DHCP address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

Step 5 (Optional) In the IPv6 Timeout field, enter the amount of time, in seconds, allowed for DHCP address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

Step 6 In the DHCP Relay Interface Servers area, add one or more interface-specific DHCP servers to which DHCP requests on a given interface are relayed:

a. Click **Add**. The Add DHCP Relay Server dialog box appears.

b. From the Interface drop-down list, choose the interface connected to the DHCP clients. Note that you do not specify the egress interface for the requests, as for a Global DHCP Server; instead, the ASA uses the routing table to determine the egress interface.

c. In the Server to... field, enter the IPv4 address of the DHCP server, and click **Add>>**. The server is added to the right-hand list. Add up to 4 servers, if available out of the overall maximum. IPv6 is not supported for interface-specific servers.

d. Click **OK**.

The newly added interface DHCP relay server(s) appear in the DHCP Relay Interface Servers list.

Step 7 To configure all interfaces as trusted interfaces, check the **Set dhcp relay information as trusted on all interfaces** check box.You can alternatively trust individual interfaces (see Step 2).

Step 8 Click **Apply** to save your settings.

Additional References

For additional information related to implementing DHCPv6, see the following section:

- RFCs, page 17-9

RFCs

RFC	Title
2132	DHCP Options and BOOTP Vendor Extensions
2462	IPv6 Stateless Address Autoconfiguration
5510	DHCP for IPv6

Monitoring DHCP Services

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 ip address dhcp lease proxy** command, then click **Send**.	Shows proxy entries in the IPL table.
Tools > Command Line Interface Enter the **show ip address dhcp lease summary** command, then click **Send**.	Shows summary for the entry.
Tools > Command Line Interface Enter the **show ip address dhcp lease server** command, then click **Send**.	Shows server entries in the IPL table.

Path	Purpose
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 addresses.
Monitoring > Interfaces > DHCP > DHCP Statistics	Shows DHCP message types, counters, values, directions, messages received, and messages sent.

Feature History for DHCP Services

Table 17-2 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 17-2 Feature History for DHCP Services

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.
DHCP relay servers per interface (IPv4 only)	9.1(2)	You can now configure DHCP relay servers per-interface, so requests that enter a given interface are relayed only to servers specified for that interface. IPv6 is not supported for per-interface DHCP relay. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP lease information	9.1(4)	You can now view DHCP Client Lease information We introduced the following screen: Monitoring > Interfaces > DHCP> DHCP Lease Information.

Table 17-2 Feature History for DHCP Services

Feature Name	Releases	Description
DHCP trusted interfaces	9.1(2)	You can now configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP rebind function	9.1(4)	During the DHCP rebind phase, the client now attempts to rebind to other DHCP servers in the tunnel group list. Prior to this release, the client did not rebind to an alternate server, when the DHCP lease fails to renew. There is no change to the ASDM.
DHCP lease information	9.1(4)	You can now view DHCP Client Lease information We introduced the following screen: Monitoring > Interfaces > DHCP> DHCP Lease Information.



Configuring Dynamic DNS

This chapter describes how to configure DDNS update methods and includes the following topics:

- Information About DDNS, page 18-1
- Licensing Requirements for DDNS, page 18-2
- Guidelines and Limitations, page 18-2
- Configuring Dynamic DNS, page 18-3
- DDNS Monitoring, page 18-7
- Feature History for DDNS, page 18-7

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 16-8. To configure DHCP, see the "Configuring the DHCP Server" section on page 17-4.

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

	Command	Purpose
Step 1	**ddns update method** *name* **Example:** `ciscoasa(config)# ddns update method ddns-2`	Creates a DDNS update method ddns-2 that dynamically updates DNS resource records (RRs).
Step 2	**ddns both** **Example:** `ciscoasa(DDNS-update-method)# ddns both`	Specifies that the client updates both the DNS A and PTR resource records (RRs).
Step 3	**interface** *mapped_name* **Example:** `ciscoasa(DDNS-update-method)# interface Ethernet0`	Configures an interface Ethernet 0and enters interface configuration mode.

Step 4	**ddns update** [*method-name* | **hostname** *hostname*]
	
	Example:
	`ciscoasa(config-if)# ddns update ddns-2`
	`ciscoasa(config-if)# ddns update hostname asa.example.com`
		Associates the the DDNS method ddns-2 with the Ethernet0 interface and an update hostname.

Step 5	**dhcp-client update dns** [**server** {**both** | **none**}] **Example:** `ciscoasa(config)# dhcp-client update dns server none`	Configures the DHCP client to request that the DHCP server perform no updates.
Step 6	**ip address dhcp** **Example:** `ciscoasa(if-config)# ip address dhcp`	Uses DHCP to obtain an IP address for the interface.
Step 7	**dhcpd update dns** [**both**] [**override**] [**interface** *srv_ifc_name*] **Example:** `ciscoasa(if-config)# dhcpd update dns both override`	Configures DHCP server to override the client update requests.

	Command	Purpose
Step 1	**interface** *mapped_name* **Example:** `ciscoasa(config)# interface Ethernet0`	Configures an interface Ethernet 0.
Step 2	**dhcp-client update dns** [**server** {**both** | **none**}] **Example:** `ciscoasa(config-if)# dhcp-client update dns both`	DHCP client requests that the DHCP server update both the DNS A and PTR resource records.
Step 3	**ddns update** [*method-name* | **hostname** *hostname*] **Example:** `ciscoasa(config-if)# ddns update hostname asa`	Configures the DHCP client on interface Ethernet 0.
Step 4	**dhcpd update dns** [**both**] [**override**] [**interface** *srv_ifc_name*] **Example:** `ciscoasa(config-if)# dhcpd update dns`	Configures DHCP server to perform DDNS updates.
Step 5	**dhcpd domain** *domain_name* [**interface** *if_name*] **Example:** `ciscoasa(config-if)# dhcpd domain example.com`	Defines the DNS domain name for DHCP clients.

	Command	Purpose
Step 1	**ddns update method** *name* **Example:** `ciscoasa(config)# ddns update method ddns-2`	Creates a DDNS update method ddns-2 that dynamically updates DNS resource records (RRs).
Step 2	**ddns** [**both**] **Example:** `ciscoasa(DDNS-update-method)# ddns`	Specifies a dynamic DNS (DDNS) update method.
Step 3	**interface** *mapped_name* **Example:** `ciscoasa(DDNS-update-method)# interface Ethernet0`	Configures an interface Ethernet 0.

	Command	Purpose
Step 4	**dhcp-client update dns** [**server** {**both** | **none**}] **Example:** ciscoasa(config-if)# dhcp-client update dns	Configures the update parameters that the DHCP client passes to the DHCP server.
Step 5	**ddns update** [*method-name* | **hostname** *hostname*] **Example:** ciscoasa(config-if)# ddns update ddns-2 ciscoasa(config-if)# ddns update hostname asa	Associates the the DDNS method ddns-2 with the Ethernet0 interface and an update hostname.
Step 6	**dhcpd update dns** [**both**] [**override**] [**interface** *srv_ifc_name*] **Example:** ciscoasa(if-config)# dhcpd update dns	Configures DHCP server to perform DDNS updates.
Step 7	**dhcpd domain** *domain_name* [**interface** *if_name*] **Example:** ciscoasa(config-if)# dhcpd domain example.com	Defines the DNS domain name for DHCP clients.

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 **show running-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 18-1 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 18-1 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.



CHAPTER 19

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 19-1
- Guidelines and Limitations, page 19-1
- Licensing Requirements for WCCP, page 19-2
- Adding or Editing WCCP Service Groups, page 19-3
- Configuring Packet Redirection, page 19-4
- WCCP Monitoring, page 19-4
- Feature History for WCCP, page 19-5

Information About WCCP

Web Cache Communication Protocol (WCCP) is a content routing protcol that allows utilization of Cisco Cache Engines (or other caches running WCCP) to localize web traffic patterns in the network, enabling content requests to be fulfilled locally. 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.

ASA Implementation of WCCP

In the ASA implementation of WCCP, the protocol interacts with other configurable features according to the following:

- AAA for network access will not work in combination with WCCP.
- An inbound access rule always takes higher priority over WCCP. For example, if an ACL does not permit a client to communicate with a server, then traffic is not redirected to a cache engine.
- 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

Does not support IPv6 traffic for redirection.

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 ACLs that include a user, user group, or a fully qualified domain name object.

Licensing Requirements for WCCP

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 19-3.

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 ACL that controls traffic redirected to this service group.

b. Choose the predefined ACL 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 19-4.

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 19-4.

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 19-3.

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.
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 19-1 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 19-1 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**



PART 5

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 20-1
- Licensing Requirements for Objects, page 20-1
- Configuring Objects, page 20-2
- Monitoring Objects, page 20-16
- Feature History for Objects, page 20-17

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 20-2
- Configuring Service Objects and Service Groups, page 20-4
- Configuring Local User Groups, page 20-7
- Configuring Security Group Object Groups, page 20-8
- Configuring Regular Expressions, page 20-10
- Configuring Time Ranges, page 20-15

Note For other objects not included in this chapter, see the following chapters:

- Local Users—See Chapter 33, "Configuring the Local Database for AAA."
- Class Maps—See Chapter 9, "Getting Started with Application Layer Protocol Inspection," in the firewall configuration guide.
- Inspect Maps—See Chapter 9, "Getting Started with Application Layer Protocol Inspection," in the firewall configuration guide.
- TCP Maps—See the Configuring Connection Settings section in the firewall configuration guide.

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 20-3
- Configuring a Network Object Group, page 20-3

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 4, "Configuring Network Object NAT," in the firewall configuration guide 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 4, "Configuring Network Object NAT," in the firewall configuration guide 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 20-5
- Configuring a Service Group, page 20-5
- Configuring a TCP or UDP Port Service Group, page 20-6

- Configuring an ICMP Group, page 20-6
- Configuring an ICMP Group, page 20-6

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 38, “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 ACLs 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.

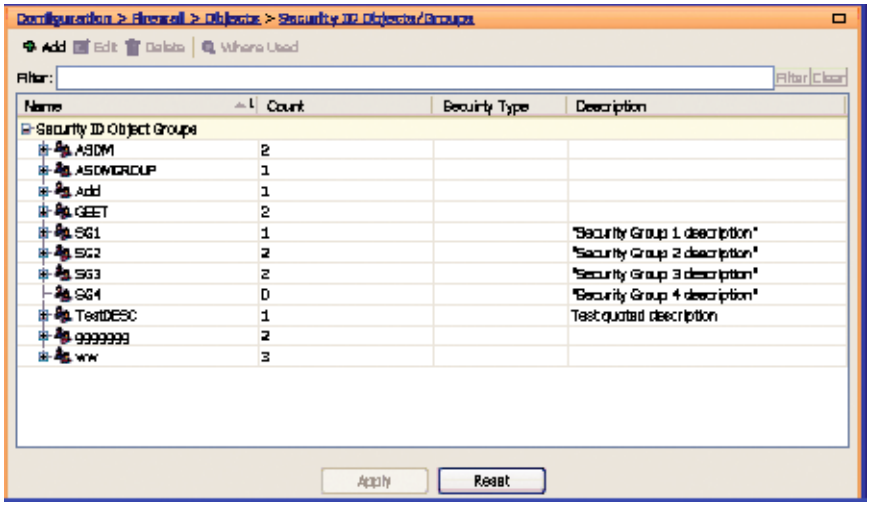
You can use the security object groups you create on the ASA to control access to network resources. You can use the security object group as part of an access group or service policy.

Prerequisites

See Chapter 39, “Configuring the ASA to Integrate with Cisco TrustSec,” to enable TrustSec.

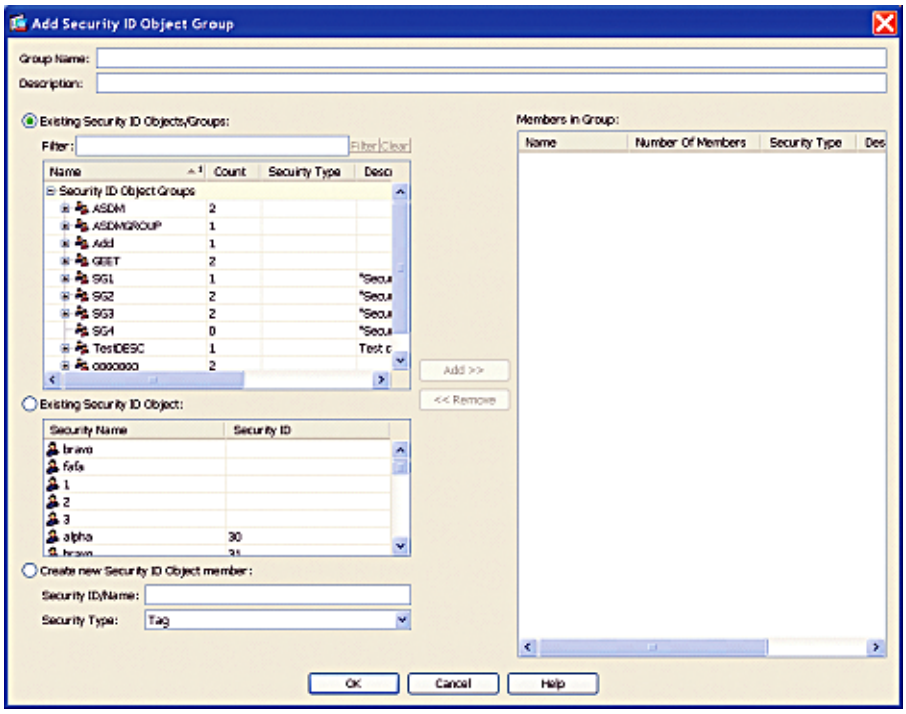
Detailed Steps

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 ACL 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 20-10
- Creating a Regular Expression Class Map, page 20-14

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 20-1 lists the metacharacters that have special meanings.

Table 20-1 regex Metacharacters

Character	Description	Notes
.	Dot	Matches any single character. For example, **d.g** matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.
(*exp*)	Subexpression	A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, **d(o|a)g** matches dog and dag, but **do|ag** matches do and ag. A subexpression can also be used with repeat quantifiers to differentiate the characters meant for repetition. For example, **ab(xy){3}z** matches abxyxyxyz.
|	Alternation	Matches either expression it separates. For example, **dog|cat** matches dog or cat.
?	Question mark	A quantifier that indicates that there are 0 or 1 of the previous expression. For example, **lo?se** matches lse or lose. **Note** You must enter **Ctrl+V** and then the question mark or else the help function is invoked.
*	Asterisk	A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, **lo*se** matches lse, lose, loose, 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.

*	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.

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 20-1 on page 20-11 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 20-1.

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 20-10.

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 ACL 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 20-2 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 20-2 Feature History for Object Groups

Feature Name	Platform Releases	Feature Information
Object groups	7.0(1)	Object groups simplify ACL creation and maintenance.
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**.
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 ACLs (also known as access control lists), and it includes the following sections:

- Information About the ACL Manager, page 21-1
- Licensing Requirements for the ACL Manager, page 21-1
- Guidelines and Limitations, page 21-2
- Adding ACLs and ACEs, page 21-2
- Feature History for the ACL Manager, page 21-5

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 21-2.

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 4-16.

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 ACL:

- Enter the ACL name in uppercase letters so that the name is easy to see in the configuration. You might want to name the ACL 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 48-11. 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 control 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 13-15.

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 41-3.

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 4-16.

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 21-1 lists the release history for this feature.

Table 21-1 Feature History for Extended ACLs

Feature Name	Releases	Feature Information
Extended ACLs	7.0(1)	ACLs 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 22-1
- Licensing Requirements for Standard ACLs, page 22-1
- Guidelines and Limitations, page 22-1
- Default Settings, page 22-2
- Adding Standard ACLs, page 22-3
- Feature History for Standard ACLs, page 22-4

Information About Standard ACLs

Standard ACLs identify the destination IP 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.

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 22-2
- Firewall Mode Guidelines, page 22-2
- IPv6 Guidelines, page 22-2
- Additional Guidelines and Limitations, page 22-2

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 22-1 lists the default settings for standard ACL parameters.

Table 22-1 Default Standard ACL Parameters

Parameters	Default
deny	The ASA denies all packets on the originating interface unless you specifically permit access. ACL 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 22-3

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 22-3
- Adding an ACE to a Standard ACL, page 22-3
- Editing an ACE in a Standard ACL, page 22-4

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 22-3.

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 22-3. For information about editing ACEs, see the "Editing an ACE in a Standard ACL" section on page 22-4

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 22-2 lists the release history for this feature.

Table 22-2 Feature History for Standard ACLs

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 23-1
- Guidelines and Limitations, page 23-1
- Default Settings, page 23-3
- Using Webtype ACLs, page 23-3
- Feature History for Webtype ACLs, page 23-6
- Feature History for Webtype ACLs, page 23-6

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 23-1
- Firewall Mode Guidelines, page 23-1
- Additional Guidelines and Limitations, page 23-2

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:

- There are two types of webtype ACLs; URL based ACLs and TCP based ACLs. URL based ACLs are used to allow or deny URLs with the format -protocol://ip-address/path, these ACLs are for filtering based on clientless features. TCP based ACLs are used to allow or deny ip-address and port.
- Permitting one type of an ACL creates an implicit deny for the other type of ACL.
- A duplicate ACE refers to ACEs with URLs that are equivalent after normalization. A duplicate ACE found during upgrade, will be removed after the upgrade.

Note URL normalization is an additional security feature that includes path normalization, case normalization and scheme normalization. URLs specified in an ACE and portal address bar are normalized before comparison; for making decisions on webvpn traffic filtering.

- If an upgrade is followed by a downgrade, duplicate ACEs will not be present in the downgraded version, if a **write memory** operation is performed after upgrade. To preserve the old configuration, you must save the running configuration to a disk, before the upgrade.
- To permit any http/https based website and all the paths within the site, www.cisco.com use the format:

```
access-list <ACL-NAME> webtype permit url http://www.cisco.com/*
```

- To permit RDP plugin protocol over clientless VPN use the format:

```
access-list <ACL-NAME> webtype permit url rdp://<host-name>/*
```

- To permit SSH plugin protocol over clientless VPN use the format:

```
access-list <ACL-NAME> webtype permit url ssh://<host-name>/*
```

- To permit telnet plugin protocol over clientless VPN use the format:

```
access-list <ACL-NAME> webtype permit url telnet://<host-name>/*
```

- To permit ica plugin protocol over clientless VPN use:

```
access-list <ACL-NAME> webtype permit url ica://<host-name>/*
```

- 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.

- Smart tunnel and ica plug-ins are not affected by an ACL with ‘permit url any’ because they match smart-tunnel:// and ica:// types.
- ‘Permit url any' will allow all the urls that have format protocol://server-ip/path and will block traffic that does not match any of the protocol://address/path such as port-forwarding; the ASA admin should explicitly set an ACE to allow connection to the required port (port 1494 in case of citrix) so that an implicit deny does not occur.

Default Settings

Table 23-1 lists the default settings for Webtype ACLs parameters.

Table 23-1 Default Webtype ACL Parameters

Parameters	Default
deny	The ASA denies all packets on the originating interface unless you specifically permit access.
log	ACL 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 23-3
- Editing Webtype ACLs and ACEs, page 23-5
- Deleting Webtype ACLs and ACEs, page 23-5

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 23-3.
- Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide 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 23-3.

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 ACLs

Table 23-2 lists the release history for this feature.

Table 23-2 Feature History for Webtype ACLs

Feature Name	Releases	Feature Information
Webtype ACLs	7.0(1)	Webtyp ACLs are ACLs that are added to a configuration that supports filtering for clientless SSL VPN. We introduced the feature.

Table 23-2 Feature History for Webtype ACLs (continued)

Feature Name	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 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
Webtype ACL enhancements	9.1(5)	• A duplicate ACE found during upgrade, will be removed after the upgrade. • If an upgrade is followed by a downgrade, duplicate ACEs will not be present in the downgraded version, if a **write memory** operation is performed after upgrade. To preserve the old configuration, you must save the running configuration to a disk, before the upgrade. **Note** A duplicate ACE refers to ACEs with URLs that are equivalent after normalization. We did not modify any ASDM screens.



PART 6

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 24-1
- How Routing Behaves Within the ASA, page 24-4
- Supported Internet Protocols for Routing, page 24-5
- Information About the Routing Table, page 24-5
- Disabling Proxy ARP Requests, page 24-11

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 24-1
- Path Determination, page 24-2
- Supported Route Types, page 24-2

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.

Supported Route Types

There are several route types that a router can use. The ASA uses the following route types:

- Static Versus Dynamic, page 24-3
- Single-Path Versus Multipath, page 24-3
- Flat Versus Hierarchical, page 24-3
- Link-State Versus Distance Vector, page 24-3

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 24-4
- Next Hop Selection Process, page 24-4

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 NAT rule 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 28-4.

- 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 27-5.

- 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 29-4.

Information About the Routing Table

This section includes the following topics:

- Displaying the Routing Table, page 24-6
- How the Routing Table Is Populated, page 24-6
- How Forwarding Decisions Are Made, page 24-8
- Dynamic Routing and Failover, page 24-9
- Dynamic Routing and Clustering, page 24-9

- Dynamic Routing in Multiple Context Mode, page 24-10

Displaying the Routing Table

To view the entries in the routing table, enter the following command:

```
hostname# show route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 10.86.194.1 to network 0.0.0.0

S    10.1.1.0 255.255.255.0 [3/0] via 10.86.194.1, outside
C    10.86.194.0 255.255.254.0 is directly connected, outside
S*   0.0.0.0 0.0.0.0 [1/0] via 10.86.194.1, outside
```

On the ASA 5505, the following route is also shown. It is the internal loopback interface, which is used by the VPN hardware client feature for individual user authentication.

```
C 127.1.0.0 255.255.0.0 is directly connected, _internal_loopback
```

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 24-1 shows the default administrative distance values for the routing protocols supported by the ASA.

Table 24-1 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 25-2.

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 24-2 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 24-2 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

EIGRP (continued)	OSPFv2 (continued)	Route Maps and Prefix Lists (continued)
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
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 ARP Requests

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 ARP requests 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 ARP requests for the interface on which you do not want them.

To disable proxy ARPs requests, enter the following commandperform the following steps:

Command	Purpose
sysopt noproxyarp *interface* **Example:** `hostname(config)# sysopt noproxyarp exampleinterface`	Disables proxy ARP requests.

Step 1 Choose **Configuration > Device Setup > Routing > Proxy ARP/Neighbor Discovery**.

The Interface field lists the interface names. The Enabled field shows whether or not proxy ARP/Neighbor Discovery is enabled (Yes) or disabled (No) for NAT global addresses.

Step 2 To enable proxy ARP/Neighbor Discovery for the selected interface, click **Enable**. By default, proxy ARP/Neighbor discovery is enabled for all interfaces.

Step 3 To disable proxy ARP/Neighbor Discovery 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 25-1
- Licensing Requirements for Static and Default Routes, page 25-2
- Guidelines and Limitations, page 25-2
- Configuring Static and Default Routes, page 25-2
- Monitoring a Static or Default Route, page 25-9
- Configuration Examples for Static or Default Routes, page 25-11
- Feature History for Static and Default Routes, page 25-13

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 25-3
- Configuring a Default Static Route, page 25-7
- Configuring IPv6 Default and Static Routes, page 25-8

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, see the following sectionchoose one of the following:

- Adding or Editing a Static Route, page 25-3
- Configuring Static Route Tracking, page 25-6
- Deleting Static Routes, page 25-6

Adding or Editing a Static Route

To add or edit a static route, enter the following command:

Command	Purpose
route *if_name dest_ip mask gateway_ip* [*distance*] **Example:** `hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.1 [1]`	Enables you to add a static route. The *dest_ip* and *mask* arguments indicate the IP address for the destination network and the *gateway_ip* argument is 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. The *distance* argument 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 route takes precedence. Connected routes always take precedence over static or dynamically discovered routes.

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.

Examples

The following example shows static routes that are equal cost routes that direct traffic to three different gateways on the outside interface. The ASA distributes the traffic among the specified gateways.

```
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.1
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.2
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.3
```

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, enter the following command:

Command	Purpose
route *if_name* 0.0.0.0 0.0.0.0 *gateway_ip* [*distance* | *tunneled*] **Example:** hostname(config)# route outside 0 0 192.168.2.4 tunneled	Enables you to add a static route. The *dest_ip* and *mask* arguments indicate the IP address for the destination network and the *gateway_ip* argument is 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. The *distance* argument 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 connect 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.

Tip You can enter 0 0 instead of 0.0.0.0 0.0.0.0 for the destination network address and mask, as shown in the following example:

```
hostname(config)# route outside 0 0 192.168.1 1
```

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 configure an IPv6 default route and static routes, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 route** *if_name* **::/0** *next_hop_ipv6_addr* **Example:** `hostname(config)# ipv6 route inside 7fff::0/32 3FFE:1100:0:CC00::1`	Adds a default IPv6 route. The example routes packets for network 7fff::0/32 to a networking device on the inside interface at 3FFE:1100:0:CC00::1 The address ::/0 is the IPv6 equivalent of any.
Step 2	**ipv6 route** *if_name destination next_hop_ipv6_addr* [*admin_distance*] **Example:** `hostname(config)# ipv6 route inside 7fff::0/32 3FFE:1100:0:CC00::1 [110]`	Adds an IPv6 static route to the IPv6 routing table. The example routes packets for network 7fff::0/32 to a networking device on the inside interface at 3FFE:1100:0:CC00::1, and with an administrative distance of 110.

Note The **ipv6 route** command works the same way as the **route** command does, which is used to define IPv4 static routes.

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 configure static route tracking, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**sla monitor** *sla_id* **Example:** `hostname(config)# sla monitor sla_id`	Configures the tracked object monitoring parameters by defining the monitoring process. If you are configuring a new monitoring process, you enter sla monitor configuration mode. If you are changing the monitoring parameters for an unscheduled monitoring process that already has a type defined, you automatically enter sla protocol configuration mode.
Step 2	**type echo protocol ipIcmpEcho** *target_ip* **interface** *if_name* **Example:** `hostname(config-sla-monitor)# type echo protocol ipIcmpEcho target_ip interface if_name`	Specifies the monitoring protocol. If you are changing the monitoring parameters for an unscheduled monitoring process that already has a type defined, you automatically enter sla protocol configuration mode and cannot change this setting. The *target_ip* argument is the IP address of the network object whose availability the tracking process monitors. While this object is available, the tracking process route is installed in the routing table. When this object becomes unavailable, the tracking process removes the route and the backup route is used in its place.
Step 3	**sla monitor schedule** *sla_id* [**life** {**forever** | *seconds*}] [**start-time** {*hh*:*mm* [:*ss*] [*month day* | *day month*] | *pending* | *now* | **after** *hh*:*mm*:*ss*}] [**ageout** *seconds*] [**recurring**] **Example:** `hostname(config)# sla monitor schedule sla_id [life {forever | seconds}] [start-time {hh:mm[:ss] [month day | day month] | pending | now | after hh:mm:ss}] [ageout seconds] [recurring]`	Schedules the monitoring process. Typically, you will use the **sla monitor schedule** *sla_id* **life forever start-time now** command for the monitoring schedule, and allow the monitoring configuration to determine how often the testing occurs. However, you can schedule this monitoring process to begin in the future and to only occur at specified times.

	Command	Purpose
Step 4	**track** *track_id* **rtr** *sla_id* **reachability** **Example:** `hostname(config)# track track_id rtr sla_id reachability`	Associates a tracked static route with the SLA monitoring process. The *track_id* argument is a tracking number you assign with this command. The *sla_id* argument is the ID number of the SLA process.
Step 5	Do one of the following to define the static route to be installed in the routing table while the tracked object is reachable. These options allow you to track a static route or a default route obtained through DHCP or PPPOE.	
	route *if_name dest_ip mask gateway_ip* [*admin_distance*] **track** *track_id* **Example:** `hostname(config)# route if_name dest_ip mask gateway_ip [admin_distance] track track_id`	Tracks a static route. You cannot use the **tunneled** option with the **route** command in static route tracking.

	Example:
	`hostname(config)# interface phy_if`
	`hostname(config-if)# dhcp client route track track_id`
	`hostname(config-if)# ip address dhcp setroute`
	`hostname(config-if)# exit`
		Tracks a default route obtained through DHCP,
		
		Remember that you must use the **setroute** keyword with the **ip address dhcp** command to obtain the default route using DHCP.

	Example:
	`hostname(config)# interface phy_if`
	`hostname(config-if)# pppoe client route track track_id`
	`hostname(config-if)# ip address pppoe setroute`
	`hostname(config-if)# exit`
		Tracks a default route obtained through PPPoE.
		
		You must use the **setroute** keyword with the **ip address pppoe** command to obtain the default route using PPPoE.

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:

```
hostname(config)# route inside 10.1.1.0 255.255.255.0 10.1.2.45 1
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.1
```

```
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.2
hostname(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.3
hostname(config)# route outside 0 0 192.168.2.4 tunneled
```

Unencrypted traffic received by the ASA for which there is no static or learned route is distributed among the gateways with the IP addresses 192.168.2.1, 192.168.2.2, and 192.168.2.3. Encrypted traffic received by the ASA for which there is no static or learned route is passed to the gateway with the IP address 192.168.2.4.

The following example creates a static route that sends all traffic destined for 10.1.1.0/24 to the router (10.1.2.45) connected to the inside interface:

```
hostname(config)# route inside 10.1.1.0 255.255.255.0 10.1.2.45 1
```

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 25-1 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 25-1 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 **route** command. 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 26-1
- Licensing Requirements for Route Maps, page 26-3
- Guidelines and Limitations, page 26-3
- Defining a Route Map, page 26-4
- Customizing a Route Map, page 26-5
- Configuration Example for Route Maps, page 26-9
- Feature History for Route Maps, page 26-10

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 ACL 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 26-4). 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 26-2
- Match and Set Clause Values, page 26-2

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 ACL 101 or ACL 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

Supports IPv6.

Additional Guidelines

Route maps do not support ACLs 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.

To define a route map, enter the following command:

Command	Purpose
route-map *name* {**permit** | **deny**} [*sequence_number*] **Example:** hostname(config)# route-map *name* {permit} [*12*]	Creates the route map entry. Enters route-map configuration mode. Route map entries are read in order. You can identify the order using the *sequence_number* argument, or the ASA uses the order in which you add route map entries.

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 26-6
- Configuring Prefix Lists, page 26-7
- Configuring Prefix Rules, page 26-8
- Configuring the Metric Values for a Route Action, page 26-8

Defining a Route to Match a Specific Destination Address

To define a route to match a specified destination address, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**route-map** *name* {**permit** | **deny**} [*sequence_number*] **Example:** hostname(config)# route-map *name* {permit} [*12*]	Creates the route map entry. Enters route-map configuration mode. Route map entries are read in order. You can identify the order using the *sequence_number* option, or the ASA uses the order in which you add route map entries.
Step 2	Enter one of the following **match** commands to match routes to a specified destination address:	
	match ip address *acl_id* [*acl_id*] [...] [prefix-list] **Example:** hostname(config-route-map)# match ip address *acl_id* [*acl_id*] [...]	Matches any routes that have a destination network that matches a standard ACL or prefix list. If you specify more than one ACL, then the route can match any of the ACLs.
	match metric *metric_value* **Example:** hostname(config-route-map)# match metric *200*	Matches any routes that have a specified metric. The *metric_value* can range from 0 to 4294967295.
	match ip next-hop *acl_id* [*acl_id*] [...] **Example:** hostname(config-route-map)# match ip next-hop *acl_id* [*acl_id*] [...]	Matches any routes that have a next hop router address that matches a standard ACL. If you specify more than one ACL, then the route can match any of the ACLs.
	match interface *if_name* **Example:** hostname(config-route-map)# match interface *if_name*	Matches any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
	match ip route-source *acl_id* [*acl_id*] [...] **Example:** hostname(config-route-map)# match ip route-source *acl_id* [*acl_id*] [...]	Matches any routes that have been advertised by routers that match a standard ACL. If you specify more than one ACL, then the route can match any of the ACLs.
	match route-type {**internal** | **external** [**type-1** | **type-2**]} **Example:** hostname(config-route-map)# match route-type internal type-1	Matches the route type.

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

If a route matches the **match** commands, then the following **set** commands determine the action to perform on the route before redistributing it.

To configure the metric value for a route action, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**route-map** *name* {**permit** | **deny**} [*sequence_number*] **Example:** `hostname(config)# route-map name {permit} [12]`	Creates the route map entry. Enters route-map configuration mode. Route map entries are read in order. You can identify the order using the *sequence_number* argument, or the ASA uses the order in which you add route map entries.
Step 2	To set a metric for the route map, enter one or more of the following **set** commands:	
	set metric *metric_value* **Example:** `hostname(config-route-map)# set metric 200`	Sets the metric value. The *metric_value* argument can range from 0 to 294967295.
	set metric-type {**type-1** | **type-2**} **Example:** `hostname(config-route-map)# set metric-type type-2`	Sets the metric type. The *metric-type* argument can be type-1 or type-2.

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.

The ASA redistributes these routes as external LSAs with a metric of 5 and a metric type of Type 1.

```
hostname(config)# route-map 1-to-2 permit
hostname(config-route-map)# match metric 1
hostname(config-route-map)# set metric 5
hostname(config-route-map)# set metric-type type-1
```

The following example shows how to redistribute the 10.1.1.0 static route into eigrp process 1 with the configured metric value:

```
hostname(config)# route outside 10.1.1.0 255.255.255.0 192.168.1.1
hostname(config-route-map)# access-list mymap2 line 1 permit 10.1.1.0 255.255.255.0
hostname(config-route-map)# route-map mymap2 permit 10
hostname(config-route-map)# match ip address mymap2
hostname(config-route-map)# router eigrp 1
hostname(config)# redistribute static metric 250 250 1 1 1 route-map mymap2
```

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 26-1 Feature History for Route Maps

Feature Name	Platform Releases	Feature Information
Route maps	7.0(1)	We introduced this feature. We introduced the following command: **route-map**. 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.
Support for Stateful Failover of dynamic routing protocols (EIGRP, OSPF, and RIP) and debugging of general routing-related operations	8.4(1)	We introduced the following commands: **debug route**, **show debug route**. We modified the following command: **show route**.
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 27-1
- Licensing Requirements for OSPF, page 27-3
- Guidelines and Limitations, page 27-3
- Configuring OSPFv2, page 27-5
- Customizing OSPFv2, page 27-7
- Configuring OSPFv3, page 27-29
- Removing the OSPF Configuration, page 27-60
- Configuration Example for OSPFv2, page 27-60
- Configuration Examples for OSPFv3, page 27-63
- Monitoring OSPF, page 27-65
- Additional References, page 27-69
- Feature History for OSPF, page 27-69

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 24-9.

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

OSPFv2 and OSPFv3 support 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 26-4. Then you generate a default route. For more information, see the "Configuring Static and Default Routes" section on page 25-2.

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 27-7.

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPF routing process and enters router configuration mode for this OSPF process. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes. If there is only one OSPF process enabled on the ASA, then that process is selected by default. You cannot change the OSPF process ID when editing an existing area.

Step 2	**network** *ip_address mask* **area** *area_id*
	
	Example:
	`hostname(config)# router ospf 2`
	`hostname(config-rtr)# network 10.0.0.0 255.0.0.0 area 0`
		Defines the IP addresses on which OSPF runs and the area ID for that interface.
		
		When adding a new area, enter the area ID. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295. You cannot change the area ID when editing an existing area.

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 27-19 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 27-13 for more information.

Customizing OSPFv2

This section explains how to customize the OSPFv2 processes and includes the following topics:

- Redistributing Routes Into OSPFv2, page 27-7
- Configuring Route Summarization When Redistributing Routes Into OSPFv2, page 27-11
- Configuring Route Summarization Between OSPFv2 Areas, page 27-13
- Configuring OSPFv2 Interface Parameters, page 27-14
- Configuring OSPFv2 Area Parameters, page 27-19
- Configuring an OSPFv2 NSSA, page 27-20
- Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3), page 27-22
- Defining Static OSPFv2 Neighbors, page 27-24
- Configuring Route Calculation Timers, page 27-25
- Logging Neighbors Going Up or Down, page 27-26
- Configuring Filtering in OSPF, page 27-27
- Configuring a Virtual Link in OSPF, page 27-28

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 25-2, and then define a route map according to the "Defining a Route Map" section on page 26-4.

To redistribute static, connected, RIP, or OSPFv2 routes into an OSPFv2 process, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	Do one of the following to redistribute the selected route type into the OSPF routing process:	

Command	Purpose
redistribute connected [[**metric** *metric-value*] [**metric-type** {**type-1** | **type-2**}] [**tag** *tag_value*] [**subnets**] [**route-map** *map_name*] **Example:** `hostname(config)# redistribute connected 5 type-1 route-map-practice`	Redistributes connected routes into the OSPF routing process.
redistribute static [**metric** *metric-value*] [**metric-type** {**type-1** | **type-2**}] [**tag** *tag_value*] [**subnets**] [**route-map** *map_name* **Example:** `hostname(config)# redistribute static 5 type-1 route-map-practice`	Redistributes static routes into the OSPF routing process.

redistribute ospf ***pid*** [**match {internal** | **external** [**1** | **2]** | **nssa-external** [**1** | **2**]}] [**metric** *metric-value*] [**metric-type** {**type-1** | **type-2**}] [**tag** *tag_value*] [**subnets**] [**route-map** *map_name*]
 	
Example:
`hostname(config)# route-map 1-to-2 permit`
`hostname(config-route-map)# match metric 1`
`hostname(config-route-map)# set metric 5`
`hostname(config-route-map)# set metric-type type-1`
`hostname(config-route-map)# router ospf 2`
`hostname(config-rtr)# redistribute ospf 1 route-map 1-to-2`
	Allows you to redistribute routes from an OSPF routing process into another OSPF routing process.
	
	You can either use the **match** options in this command to match and set route properties, or you can use a route map. The **subnets** option does not have equivalents in the **route-map** command. If you use both a route map and **match** options in the **redistribute** command, then they must match.
	
	The example shows route redistribution from OSPF process 1 into OSPF process 2 by matching routes with a metric equal to 1. The ASA redistributes these routes as external LSAs with a metric of 5 and a metric type of Type 1.

Command	Purpose

redistribute rip [**metric** *metric-value*] [**metric-type** {**type-1** | **type-2**}] [**tag** *tag_value*] [**subnets**] [**route-map** *map_name*]
 	
Example:
`hostname(config)# redistribute rip 5`
`hostname(config-route-map)# match metric 1`
`hostname(config-route-map)# set metric 5`
`hostname(config-route-map)# set metric-type type-1`
`hostname(config-rtr)# redistribute ospf 1 route-map 1-to-2`
	Allows you to redistribute routes from a RIP routing process into the OSPF routing process.

redistribute eigrp ***as-num*** [**metric** *metric-value*] [**metric-type** {**type-1** | **type-2**}] [**tag** *tag_value*] [**subnets**] [**route-map** *map_name*]
 	
Example:
`hostname(config)# redistribute eigrp 2`
`hostname(config-route-map)# match metric 1`
`hostname(config-route-map)# set metric 5`
`hostname(config-route-map)# set metric-type type-1`
`hostname(config-rtr)# redistribute ospf 1 route-map 1-to-2`
	Allows you to redistribute routes from an EIGRP routing process into the OSPF routing process.

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 26-4.

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 27-11
- Adding or Editing an OSPF Summary Address, page 27-12

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 1`	Creates an OSPF routing process and enters router configuration mode for this OSPF process. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.

Step 2	`summary-address` *ip_address mask* `[`**not-advertise**`] [tag` *tag*`]`
	
	Example:
	`hostname(config)# router ospf 1`
	`hostname(config-rtr)# summary-address 10.1.0.0 255.255.0.0`
		Sets the summary address.
		
		In this example, the summary address 10.1.0.0 includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the 10.1.0.0 address is advertised in an external link-state advertisement.

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 1`	Creates an OSPF routing process and enters router configuration mode for this OSPF process. The *process_id* argument is an internally used identifier for this routing process. It can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.

Step 2	`area area-id range ip-address mask [advertise | not-advertise]`
	
	Example:
	`hostname(config)# router ospf 1`
	`hostname(config-rtr)# area 17 range 12.1.0.0 255.255.0.0`
		Sets the address range.
		
		In this example, the address range is set between OSPF areas.

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: **ospf hello-interval**, **ospf dead-interval**, and **ospf authentication-key**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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.

Step 2	**network** *ip_address mask* **area** *area_id*
	
	Example:
	`hostname(config)# router ospf 2`
	`hostname(config-rtr)# network 10.0.0.0 255.0.0.0 area 0`
		Defines the IP addresses on which OSPF runs and the area ID for that interface.

Step 3	**interface** *interface_name* **Example:** `hostname(config)# interface my_interface`	Allows you to enter interface configuration mode.
Step 4	Do one of the following to configure optional OSPF interface parameters:	
	ospf authentication [**message-digest** | **null**] **Example:** `hostname(config-interface)# ospf authentication message-digest`	Specifies the authentication type for an interface.

Command	Purpose
ospf authentication-key *key* **Example:** `hostname(config-interface)# ospf authentication-key cisco`	Allows you to assign a password to be used by neighboring OSPF routers on a network segment that is using the OSPF simple password authentication. The *key* argument can be any continuous string of characters up to 8 bytes in length. The password created by this command is used as a key that is inserted directly into the OSPF header when the ASA software originates routing protocol packets. A separate password can be assigned to each network on a per-interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.
ospf cost *cost* **Example:** `hostname(config-interface)# ospf cost 20`	Allows you to explicitly specify the cost of sending a packet on an OSPF interface. The *cost* is an integer from 1 to 65535. In this example, the cost is set to 20.
ospf dead-interval *seconds* **Example:** `hostname(config-interface)# ospf dead-interval 40`	Allows you to set the number of seconds that a device must wait before it declares a neighbor OSPF router down because it has not received a hello packet. The value must be the same for all nodes on the network. In this example, the dead interval is set to 40.
ospf hello-interval *seconds* **Example:** `hostname(config-interface)# ospf hello-interval 10`	Allows you to specify the length of time between the hello packets that the ASA sends on an OSPF interface. The value must be the same for all nodes on the network. In this example, the hello interval is set to 10.
ospf message-digest-key *key_id* **md5** *key* **Example:** `hostname(config-interface)# ospf message-digest-key 1 md5 cisco`	Enables OSPF MD5 authentication. The following argument values can be set: • *key_id*—An identifier in the range from 1 to 255. • *key*—An alphanumeric password of up to 16 bytes. Usually, one key per interface is used to generate authentication information when sending packets and to authenticate incoming packets. The same key identifier on the neighbor router must have the same key value. We recommend that you not keep more than one key per interface. Every time you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. Removing the old key also reduces overhead during rollover.
ospf priority *number_value* **Example:** `hostname(config-interface)# ospf priority 20`	Allows you to set the priority to help determine the OSPF designated router for a network. The *number_value* argument ranges from 0 to 255. In this example, the priority number value is set to 20.

Command	Purpose
ospf retransmit-interval *seconds* **Example:** `hostname(config-interface)# ospf retransmit-interval seconds`	Allows you to specify the number of seconds between LSA retransmissions for adjacencies belonging to an OSPF interface. The value for *seconds* must be greater than the expected round-trip delay between any two routers on the attached network. The range is from 1 to 65535 seconds. The default value is 5 seconds. In this example, the retransmit-interval value is set to 15.
ospf transmit-delay *seconds* **Example:** `hostname(config-interface)# ospf transmit-delay 5`	Sets the estimated number of seconds required to send a link-state update packet on an OSPF interface. The *seconds* value ranges from 1 to 65535 seconds. The default value is 1 second. In this example, the transmit-delay is 5 seconds.
ospf network point-to-point non-broadcast **Example:** `hostname(config-interface)# ospf network point-to-point non-broadcast`	Specifies the interface as a point-to-point, non-broadcast network. When you designate an interface as point-to-point and non-broadcast, you must manually define the OSPF neighbor; dynamic neighbor discovery is not possible. See the "Defining Static OSPFv2 Neighbors" section on page 27-24 for more information. Additionally, you can only define one OSPF neighbor on that interface.

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 27-19 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 27-24 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 25-2 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 further reduce the number of LSAs sent into a stub area, you can use the **no-summary** keyword of the **area stub** command on the ABR to prevent it from sending a summary link advertisement (LSA Type 3) into the stub area.

To specify OSPFv2 area parameters for your network, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	Do one of the following to configure optional OSPF area parameters:	
	area *area-id* **authentication** **Example:** `hostname(config-rtr)# area 0 authentication`	Enables authentication for an OSPF area.
	area *area-id* **authentication message-digest** **Example:** `hostname(config-rtr)# area 0 authentication message-digest`	Enables MD5 authentication for an OSPF area.

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPF routing process and enters router configuration mode for the OSPF routing process that you want to redistribute. The *process_id* argument is an internally used identifier for this routing process. It can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	Do one of the following to configure optional OSPF NSSA parameters:	
	area *area-id* **nssa** [**no-redistribution**] [**default-information-originate**] **Example:** `hostname(config-rtr)# area 0 nssa`	Defines an NSSA area.
	`summary-address` *ip_address mask* [**not-advertise**] [`tag` *tag*] **Example:** `hostname(config-rtr)# summary-address 10.1.0.0 255.255.0.0`	Sets the summary address and helps reduce the size of the routing table. Using this command for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. In this example, the summary address 10.1.0.0 includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the 10.1.0.0 address is advertised in an external link-state advertisement.

Note OSPF does not support summary-address 0.0.0.0 0.0.0.0.

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 and OSPFv3, enter the following command:

Command	Purpose

router-id cluster-pool hostname | **A.B.C.D** *ip_pool*
 	
Example:
`hostname(config)# ip local pool rpool 1.1.1.1-1.1.1.4`
`hostname(config)# router ospf 1`
`hostname(config-rtr)# router-id cluster-pool rpool`
`hostname(config-rtr)# network 17.5.0.0 255.255.0.0 area 1`
`hostname(config-rtr)# log-adj-changes`
	Specifies the router ID cluster pool for Layer 3 clustering.
	
	The **cluster-pool** keyword enables configuration of an IP address pool when Layer 3 clustering is configured. The **hostname** | **A.B.C.D.** keyword specifies the OSPF router ID for this OSPF process. The *ip_pool* argument specifies the name of the IP address pool.
	
	Note If you are using clustering, then you do not need to specify an IP address pool for the router ID. If you do not configure an IP address pool, then the ASA uses the automatically generated router ID.

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 25, "Configuring Static and Default Routes," for more information about creating static routes.

To define a static OSPFv2 neighbor, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	**neighbor** *addr* [**interface** *if_name*] **Example:** `hostname(config-rtr)# neighbor 255.255.0.0 [interface my_interface]`	Defines the OSPFv2 neighborhood. The *addr* argument is the IP address of the OSPFv2 neighbor. The *if_name* argument is the interface used to communicate with the neighbor. If the OSPF v2neighbor is not on the same network as any of the directly connected interfaces, you must specify the interface.

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	**timers spf** *spf-delay spf-holdtime* **Example:** `hostname(config-rtr)# timers spf 10 120`	Configures the route calculation times. The *spf-delay* argument is the delay time (in seconds) between when OSPF receives a topology change and when it starts an SPF calculation. It can be an integer from 0 to 65535. The default time is 5 seconds. A value of 0 means that there is no delay; that is, the SPF calculation is started immediately. The *spf-holdtime* argument is the minimum time (in seconds) between two consecutive SPF calculations. It can be an integer from 0 to 65535. The default time is 10 seconds. A value of 0 means that there is no delay; that is, two SPF calculations can be performed, one immediately after the other.

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.

Configure the **log-adj-changes** command if you want to know about OSPFv2 neighbors going up or down without turning on the **debug ospf adjacency** command. The **log-adj-changes** command provides a higher level view of the peer relationship with less output. Configure the **log-adj-changes detail** command if you want to see messages for each state change.

To log OSPFv2 neighbors going up or down, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config)# router ospf 2`	Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process. The *process_id* argument is an internally used identifier for this routing process and can be any positive integer. This ID does not have to match the ID on any other device; it is for internal use only. You can use a maximum of two processes.
Step 2	**log-adj-changes** [**detail**] **Example:** `hostname(config-rtr)# log-adj-changes [detail]`	Configures logging for neighbors going up or down.

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 26-7 and the “Configuring Prefix Rules” section on page 26-8.

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 27-19 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 27-30
- Configuring OSPFv3 Interface Parameters, page 27-31
- Configuring OSPFv3 Router Parameters, page 27-38
- Configuring OSPFv3 Area Parameters, page 27-40
- Configuring a Virtual Link Neighbor, page 27-42
- Configuring OSPFv3 Passive Interfaces, page 27-44
- Configuring OSPFv3 Administrative Distance, page 27-45
- Configuring OSPFv3 Timers, page 27-46
- Defining Static OSPFv3 Neighbors, page 27-50
- Resetting OSPFv3 Default Parameters, page 27-52
- Sending Syslog Messages, page 27-53
- Suppressing Syslog Messages, page 27-54
- Calculating Summary Route Costs, page 27-54
- Generating a Default External Route into an OSPFv3 Routing Domain, page 27-55
- Configuring an IPv6 Summary Prefix, page 27-56
- Redistributing IPv6 Routes, page 27-57

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, enter the following command or perform the following stepsperform the following steps:

Command	Purpose
ipv6 router ospf *process-id* **Example:** `hostname(config)# ipv6 router ospf 10`	Creates an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used tag for this routing process and can be any positive integer. This tag does not have to match the tag on any other device; it is for internal use only. You can use a maximum of two processes.

	Command	Purpose
Step 1	**interface** *interface_name* **Example:** `hostname(config)# interface Gigabitethernet0/0`	Enables an interface.
Step 2	**ipv6 ospf** *process-id* **area** *area_id* **Example:** `hostname(config)# ipv6 ospf 200 area 100`	Creates the OSPFv3 routing process with the specified process ID and an area for OSPFv3 with the specified area ID.

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 27-40.

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: **ipv6 ospf hello-interval** and **ipv6 ospf dead-interval**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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 10`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used tag for this routing process and can be any positive integer. This tag does not have to match the tag on any other device; it is for internal use only. You can use a maximum of two processes.

Step 2	**ipv6 ospf area** [*area-num*] [**instance**]
	
	Example:
	`hostname(config-if)# interface GigabitEthernet3/2.200`
	`vlan 200`
	`nameif outside`
	`security-level 100`
	`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
	`ipv6 address 3001::1/64 standby 3001::8`
	`ipv6 address 6001::1/64 standby 6001::8`
	`ipv6 enable`
	`ospf priority 255`
	`ipv6 ospf cost 100`
	`ipv6 ospf 100 area 10 instance 200`
		Creates an OSPFv3 area.
		
		The *area-num* argument is the area for which authentication is to be enabled and can be either a decimal value or an IP address. The **instance** keyword specifies the area instance ID that is to be assigned to an 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 3	Do one of the following to configure OSPFv3 interface parameters:	

	ipv6 ospf cost *interface-cost*
	
	Example:
	`hostname(config-if)# interface GigabitEthernet3/2.200`
	`vlan 200`
	`nameif outside`
	`security-level 100`
	`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
	`ipv6 address 3001::1/64 standby 3001::8`
	`ipv6 address 6001::1/64 standby 6001::8`
	`ipv6 enable`
	`ospf priority 255`
	`ipv6 ospf cost 100`
	`ipv6 ospf 100 area 10 instance 200`
		Explicitly specifies the cost of sending a packet on an interface. The *interface-cost* argument specifies an unsigned integer value expressed as the link-state metric, which can range in value from 1 to 65535. The default cost is based on the bandwidth.

Command	Purpose

ipv6 ospf database-filter all out
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf database-filter all out`
	Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.

ipv6 ospf dead-interval *seconds*
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf dead-interval 60`
	Sets 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. The default is four times the interval set by the **ipv6 ospf hello-interval** command.

Command	Purpose

ipv6 ospf encryption {**ipsec spi** *spi* **esp** *encryption-algorithm* [[*key-encryption-type*] *key*] *authentication-algorithm* [[*key-encryption-type*] *key* | **null**}
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf encryption ipsec spi 1001 esp null sha1 123456789A123456789B123456789C123456789D`
	Specifies the encryption type for an interface. The **ipsec** keyword specifies the IP security protocol. The **spi** *spi* keyword-argument pair specifies the security policy index, which must be in the range of 256 to 42949667295 and entered as a decimal. The **esp** keyword specifies the encapsulating security payload. The *encryption-algorithm* argument specifies the encryption algorithm to be used with ESP. Valid values include the following:
	• aes-cdc—Enables AES-CDC encryption.
	• 3des—Enables 3DES encryption.
	• des—Enables DES encryption.
	• null—Specifies ESP with no encryption.
	The *key-encryption-type* argument can be one of the following two values:
	• 0—The key is not encrypted.
	• 7—The key is encrypted.
	The *key* argument specifies the number used in the calculation of the message digest. The number is 32 hexadecimal digits (16 bytes) long. The size of the key depends on the encryption algorithm used. Some algorithms, such as AES-CDC, allow you to choose the size of the key. The *authentication-algorithm* argument specifies the encryption authentication algorithm to be used, which can be one of the following:
	• md5—Enables message digest 5 (MD5).
	• sha1—Enables SHA-1.
	The **null** keyword overrides area encryption.
	Note If OSPFv3 encryption is enabled on an interface and a neighbor is on different area (for example, area 0), and you want the ASA to form adjacencies with that area, you must change the area on the ASA. After you have changed the area on the ASA to 0, there is a delay of two minutes before the OSPFv3 adjacency comes up.

ipv6 ospf flood-reduction
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf flood reduction`
	Specifies the flood reduction of LSAs to the interface.

Command	Purpose

ipv6 ospf hello-interval *seconds*
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf hello-interval 15`
	Specifies 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.

ipv6 ospf mtu-ignore
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf mtu-ignore`
	Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.

ipv6 ospf network {**broadcast** | **point-to-point non-broadcast**}
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf network point-to-point non-broadcast`
	Sets the OSPF network type to a type other than the default, which depends on the network type. The **point-to-point non-broadcast** keyword sets the network type to point-to-point non-broadcast. The **broadcast** keyword sets the network type to broadcast.

Command	Purpose

ipv6 ospf priority *number-value*
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf priority 4`
	Sets the router priority, which helps determine the designated router for a network. Valid values range from 0 to 255.

ipv6 ospf neighbor *ipv6-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*] [**database-filter all out**]
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf neighbor FE80::A8BB:CCFF:FE00:C01`
	Configures OSPFv3 router interconnections to non-broadcast networks.

Command	Purpose

ipv6 ospf retransmit-interval *seconds*
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf retransmit-interval 8`
	Specifies 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.

ipv6 ospf transmit-delay *seconds*
 	
Example:
`hostname(config-if)# interface GigabitEthernet3/2.200`
`vlan 200`
`nameif outside`
`security-level 100`
`ip address 10.20.200.30 255.255.255.0 standby 10.20.200.31`
`ipv6 address 3001::1/64 standby 3001::8`
`ipv6 address 6001::1/64 standby 6001::8`
`ipv6 enable`
`ospf priority 255`
`ipv6 ospf cost 100`
`ipv6 ospf 100 area 10 instance 200`
`ipv6 ospf retransmit-delay 3`
	Sets 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 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 Router Parameters

To configure OSPFv3 router parameters for IPv6, perform the following steps:

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config)# ipv6 router ospf 10`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	Do one of the following to configure optional OSPFv3 router parameters:	

	`area`
	
	Example:
	`hostname(config-rtr)# area 10`
		Configures OSPFv3 area parameters. Supported parameters include the area ID as a decimal value from 0 to 4294967295 and the area ID in the IP address format of **A.B.C.D**.

	default **Example:** `hostname(config-rtr)# default originate`	Sets a command to its default value. The **originate** parameter distributes the default route.
	default-information **Example:** `hostname(config-rtr)# default-information`	Controls distribution of default information.
	distance **Example:** `hostname(config-rtr)# distance 200`	Defines the OSPFv3 route administrative distance based on the route type. Supported parameters include the administrative distance with values from 1 to 254 and **ospf** for the OSPFv3 distance.

	`exit`
	
	Example:
	`hostname(config-rtr)# exit`
		Exits from IPv6 router configuration mode.

	ignore **Example:** `hostname(config-rtr)# ignore lsa`	Suppresses the sending of syslog messages with the **lsa** parameter when the router receives a link-state advertisement (LSA) for Type 6 Multicast OSPF (MOSPF) packets.
	log-adjacency-changes **Example:** `hostname(config-rtr)# log-adjacency-changes detail`	Configures the router to send a syslog message when an OSPFv3 neighbor goes up or down. With the **detail** parameter, all state changes are logged.

Command	Purpose
passive-interface [*interface_name*] **Example:** `hostname(config-rtr)# passive-interface inside`	Suppresses the sending and receiving of routing updates on an interface. The *interface_name* argument specifies the name of the interface on which the OSPFv3 process is running.
redistribute **Example:** `hostname(config-rtr)# redistribute ospf`	Configures the redistribution of routes from one routing domain into another according to the following parameters: • **connected**—Specifies connected routes. • **ospf**—Specifies OSPFv3 routes. • **static**—Specifies static routes.
router-id **Example:** `hostname(config-rtr)# router-id 10.1.1.1`	Creates a fixed router ID for a specified process with the following parameters: • **A.B.C.D**—Specifies the OSPF router ID in IP address format. • **cluster-pool**—Configures an IP address pool when Layer 3 clustering is configured. For more information about IP address pools used in clustering, see the "Configuring an IP Address Pool for Clustering (OSPFv2 and OSPFv3)" section on page 27-22.

summary-prefix
 	
Example:
`hostname(config-if)# ipv6 router ospf 1`
`hostname(config-router)# router-id 192.168.3.3`
`hostname(config-router)# summary-prefix FECO::/24`
`hostname(config-router)# redistribute static`
	Configures IPv6 address summaries with valid values from 0 to 128. The **X:X:X:X::X/** parameter specifies the IPv6 prefix.

timers
 	
Example:
`hostname(config)# ipv6 router ospf 10`
`hostname(config-rtr)# timers throttle spf 6000 12000 14000`
	Adjusts routing timers. The routing timer parameters are the following:
	• **lsa**—Specifies OSPFv3 LSA timers.
	• **pacing**—Specifies OSPFv3 pacing timers.
	• **throttle**—Specifies OSPFv3 throttle timers.

Configuring OSPFv3 Area Parameters

To configure OSPFv3 area parameters, perform the following steps:

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	`Do one of the following to configure optional OSPFv3 area parameters:`	
	area *area-id* **default-cost** *cost* **Example:** `hostname(config-rtr)# area 1 default-cost nssa`	Sets the summary default cost of an NSSA area or a stub area.
	area *area-id* **range** *ipv6-prefix/prefix-length* [**advertise** | **not advertise**] [**cost** *cost*] **Example:** `hostname(config-rtr)# area 1 range FE01:1::1/64`	Summarizes routes that match the address and mask for border routers only. The *area-id* argument identifies the area for which routes are to be summarized. The value can be specified as a decimal or an IPv6 prefix. The *ipv6-prefix* argument specifies the IPv6 prefix. The *prefix-length* argument specifies the prefix length. The **advertise** keyword sets the address range status to advertised and generates a Type 3 summary LSA. The **not-advertise** keyword sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks. The **cost** *cost* keyword-argument pair specifies 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.
	area *area-id* **nssa** **Example:** `hostname(config-rtr)# area 1 nssa`	Specifies an NSSA area.

Command	Purpose
area *area-id* **stub** **Example:** `hostname(config-rtr)# area 1 stub`	Specifies a stub area.
area *area-id* **virtual-link** *router-id* [**hello-interval** *seconds*] [**retransmit-interval** *seconds*] [**transmit-delay** *seconds*] [**dead-interval** *seconds*] [**ttl-security hops** *hop-count*] **Example:** `hostname(config-rtr)# area 1 virtual-link 192.168.255.1 hello-interval 5`	Defines a virtual link and its parameters. The *area-id* argument identifies the area for which routes are to be summarized. The **virtual link** keyword specifies the creation of a virtual link neighbor. The *router-id* argument specifies the router ID that is associated with the virtual link neighbor. Enter the **show ip ospf** or **show ipv6 display** command to display the router ID. There is no default value. The **hello-interval** keyword specifies the time in seconds between the hello packets that are sent on an interface. 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. The **retransmit-interval** *seconds* keyword-argument pair specifies the time in seconds between LSA retransmissions for adjacencies that belong to the interface. 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. The **transmit-delay** *seconds* keyword-argument pair specifies the estimated time in seconds that is required to send a link-state update packet on the interface. 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. The **dead-interval** *seconds* keyword-argument pair specifies the time in seconds that hello packets are not seen before a neighbor indicates that the router is down. 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. The **ttl-security hops** keyword configures the time-to-live (TTL) security on a virtual link. The *hop-count* argument value can range from 1 to 254.

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 27-42.

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:

	Command	Purpose
Step 1	**ipv6 router ospf** *process_id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process_id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**passive-interface** [*interface_name*] **Example:** `hostname(config-rtr)# passive-interface inside`	Suppresses the sending and receiving of routing updates on an interface. The *interface_name* argument specifies the name of the interface on which the OSPFv3 process is running. If the *no interface_name* argument is specified, all of the interfaces in the OSPFv3 process *process_id* are made passive.

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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process_id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process_id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**distance** [**ospf** {**external** | **inter-area** | **intra-area**}] *distance* **Example:** `hostname(config-rtr)# distance ospf external 200`	Sets the administrative distance for OSPFv3 routes. The **ospf** keyword specifies OSPFv3 routes. The **external** keyword specifies the external Type 5 and Type 7 routes for OSPFv3. The **inter-area** keyword specifies the inter-area routes for OSPVv3. The **intra-area** keyword specifies the intra-area routes for OSPFv3. The *distance* argument specifies the administrative distance, which is an integer from 10 to 254.

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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id*	Enables an OSPFv3 routing process and enters IPv6 router configuration mode.
	Example: `hostname(config-if)# ipv6 router ospf 1`	The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**timers lsa arrival** *milliseconds*	Sets the minimum interval at which the ASA accepts the same LSA from OSPF neighbors.
	Example: `hostname(config-rtr)# timers lsa arrival 2000`	The *milliseconds* argument specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6,000,000 milliseconds. The default is 1000 milliseconds.

To configure LSA flood packet pacing, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id*	Enables an OSPFv3 routing process and enters IPv6 router configuration mode.
	Example: `hostname(config-if)# ipv6 router ospf 1`	The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**timers pacing flood** *milliseconds*	Configures LSA flood packet pacing.
	Example: `hostname(config-rtr)# timers lsa flood 20`	The *milliseconds* argument 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.

To change the interval at which OSPFv3 LSAs are collected into a group and refreshed, check summed, or aged, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id*	Enables an OSPFv3 routing process and enters IPv6 router configuration mode.
	Example: `hostname(config-if)# ipv6 router ospf 1`	The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**timers pacing lsa-group** *seconds*	Changes the interval at which OSPFv3 LSAs are collected into a group and refreshed, checksummed, or aged.
	Example: `hostname(config-rtr)# timers pacing lsa-group 300`	The *seconds* argument specifies the number of seconds in the interval at which LSAs are grouped, refreshed, check summed, or aged. The range is from 10 to 1800 seconds. The default value is 240 seconds.

To configure LSA retransmission packet pacing, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id*	Enables an OSPFv3 routing process and enters IPv6 router configuration mode.
	Example: `hostname(config-if)# ipv6 router ospf 1`	The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**timers pacing retransmission** *milliseconds*	Configures LSA retransmission packet pacing.
	Example: `hostname(config-rtr)# timers pacing retransmission 100`	The *milliseconds* argument 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.

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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	Choose one of the following options:	

Command	Purpose
timers throttle lsa *milliseconds1 milliseconds2 milliseconds3* **Example:** `hostname(config-rtr)# timers throttle lsa 500 6000 8000`	Configures OSPFv3 LSA throttling. The *milliseconds1* argument specifies the delay in milliseconds to generate the first occurrence of the LSA. The *milliseconds2* argument specifies the maximum delay in milliseconds to originate the same LSA. The *milliseconds3* argument specifies the minimum delay in milliseconds to originate the same LSA. 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 default values for LSA throttling are the following: • For *milliseconds1*, the default value is 0 milliseconds. • For *milliseconds2* and *milliseconds3*, the default value is 5000 milliseconds.
timers throttle spf *milliseconds1 milliseconds2 milliseconds3* **Example:** `hostname(config-rtr)# timers throttle spf 5000 12000 16000`	Configures OSPFv3 SPF throttling. The *milliseconds1* argument specifies the delay in milliseconds to receive a change to the SPF calculation. The *milliseconds2* argument specifies the delay in milliseconds between the first and second SPF calculations. The *milliseconds3* argument specifies the maximum wait time in milliseconds for SPF calculations. For SPF throttling, if *milliseconds2* or *milliseconds3* is less than *milliseconds1*, then OSPFv3 automatically corrects to the *milliseconds1* value. Similarly, if *milliseconds3* is less than *milliseconds2*, then OSPFv3 automatically corrects to the *milliseconds2* value. The default values for SPF throttling are the following: • For *milliseconds1*, the default value is 5000 milliseconds. • For *milliseconds2* and *milliseconds3*, the default value is 10000 milliseconds.

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 25, “Configuring Static and Default Routes,” for more information about creating static routes.

To define a static OSPFv3 neighbor, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.

Step 2	**ipv6 ospf neighbor** *ipv6-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*] [**database-filter all out**]
	
	Example:
	`hostname(config-if)# interface ethernet0/0`
	`ipv6 ospf neighbor`
	`FE80::A8BB:CCFF:FE00:C01`
		Configures OSPFv3 router interconnections to non-broadcast networks.

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**.

Resetting OSPFv3 Default Parameters

To return an OSPFv3 parameter to its default value, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process_id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**default** [**area** | **auto-cost** | **default-information** | **default-metric** | **discard-route** | **discard-route** | **distance** | **distribute-list** | **ignore** | **log-adjacency-changes** | **maximum-paths** | **passive-interface** | **redistribute** | **router-id** | **summary-prefix** | **timers**] **Example:** `hostname(config-rtr)# default metric 5`	Returns an optional parameter to its default value. The **area** keyword specifies the OSPFv3 area parameters. The **auto-cost** keyword specifies the OSPFv3 interface cost according to bandwidth. The **default-information** keyword distributes default information. The **default-metric** keyword specifies the metric for a redistributed route. The **discard-route** keyword enables or disables the discard-route installation. The **distance** keyword specifies the administrative distance. The **distribute-list** keyword filters networks in routing updates. The **ignore** keyword ignores a specific event. The **log-adjacency-changes** keyword logs changes in the adjacency state. The **maximum-paths** keyword forwards packets over multiple paths. The **passive-interface** keyword suppresses routing updates on an interface. The **redistribute** keyword redistributes IPv6 prefixes from another routing protocol. The **router-id** keyword specifies the router ID for the specified routing process. The **summary-prefix** keyword specifies the IPv6 summary prefix. The **timers** keyword specifies the OSPFv3 timers.

Sending Syslog Messages

To configure the router to send a syslog message when an OSPFv3 neighbor goes up or down, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**log-adjacency-changes** [**detail**] **Example:** `hostname(config-rtr)# log-adjacency-changes detail`	Configures the router to send a syslog message when an OSPFv3 neighbor goes up or down. The **detail** keyword sends a syslog message for each state, not only when an OSPFv3 neighbor goes up or down.

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:

Detailed Steps

	Command	Purpose
Step 1	**router ospf** *process_id* **Example:** `hostname(config-if)# router ospf 1`	Enables an OSPFv2 routing process and enters router configuration mode. The *process_id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**ignore lsa mospf** **Example:** `hostname(config-rtr)# ignore lsa mospf`	Suppresses the sending of syslog messages when the router receives unsupported LSA Type 6 MOSPF packets.

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, enter the following commandperform the following steps:

Command	Purpose
compatible rfc1583 **Example:** `hostname (config-rtr)# compatible rfc1583`	Restores the methods that are used to calculate summary route costs according to RFC 1583.

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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**default-information originate** [**always**] **metric** *metric-value* [**metric-type** *type-value*] [**route-map** *map-name*] **Example:** `hostname(config-rtr)# default-information originate always metric 3 metric-type 2`	Generates a default external route into an OSPFv3 routing domain. The **always** keyword advertises the default route whether or not the default route exists. The **metric** *metric-value* keyword-argument pair specifies the metric used for generating the default route. If you do not specify a value using the **default-metric** command, the default value is 10. Valid metric values range from 0 to 16777214. The **metric-type** *type-value* keyword-argument pair specifies the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values can be one of the following: • 1—Type 1 external route • 2—Type 2 external route The default is the type 2 external route. The **route-map** *map-name* keyword-argument pair specifies the routing process that generates the default route if the route map is satisfied.

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:

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process_id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.

Step 2	**summary-prefix** *prefix* [**not-advertise** | **tag** *tag-value*]
	
	Example:
	`hostname(config-if)# ipv6 router ospf 1`
	`hostname(config-rtr)# router-id 192.168.3.3`
	`hostname(config-rtr)# summary-prefix FECO::/24`
	`hostname(config-rtr)# redistribute static`
		Configures an IPv6 summary prefix.
		
		The *prefix* argument is the IPv6 route prefix for the destination. The **not-advertise** keyword suppresses routes that match the specified prefix and mask pair. This keyword applies to OSPFv3 only. The **tag** *tag-value* keyword-argument pair specifies the tag value that can be used as a match value for controlling redistribution through route maps. This keyword applies to OSPFv3 only.

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.

Detailed Steps

	Command	Purpose
Step 1	**ipv6 router ospf** *process-id* **Example:** `hostname(config-if)# ipv6 router ospf 1`	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The *process-id* argument is an internally used identifier for this routing process, is locally assigned, and can be any positive integer from 1 to 65535. This ID does not have to match the ID on any other device; it is for internal administrative use only. You can use a maximum of two processes.
Step 2	**redistribute** *source-protocol* [*process-id*] [**include-connected** {[**level-1** | **level-2**}] [*as-number*] [**metric** [*metric-value* | **transparent**}] [**metric-type** *type-value*] [**match** {**external** [**1**|**2**] | **internal** | **nssa-external** [**1**|**2**]}] [**tag** *tag-value*] [**route-map** *map-tag*] **Example:** `hostname(config-rtr)# redistribute connected 5 type-1`	Redistributes IPv6 routes from one OSPFv3 process into another. The *source-protocol* argument specifies the source protocol from which routes are being redistributed, which can be static, connected, or OSPFv3. The *process-id* argument is the number that is assigned administratively when the OSPFv3 routing process is enabled. The **include-connected** keyword allows the target protocol to redistribute routes learned by the source protocol and connected prefixes on those interfaces over which the source protocol is running. The **level-1** keyword specifies that for Intermediate System-to-Intermediate System (IS-IS), Level 1 routes are redistributed into other IP routing protocols independently. The **level-1-2** keyword specifies that for IS-IS, both Level 1 and Level 2 routes are redistributed into other IP routing protocols. The **level-2** keyword specifies that for IS-IS, Level 2 routes are redistributed into other IP routing protocols independently. For the **metric** *metric-value* keyword-argument pair, when redistributing routes from one OSPFv3 process into another OSPFv3 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 OSPFv3 process, the default metric is 20 when no metric value is specified. The **metric transparent** keyword causes RIP to use the routing table metric for redistributed routes as the RIP metric. The **metric-type** *type-value* keyword-argument pair specifies the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values can be one of the following: 1 for a Type 1 external route or 2 for a Type 2 external route. If no value is specified for the **metric-type** keyword, the ASA adopts a Type 2 external route. For IS-IS, the link type can be one of the following: internal for an IS-IS metric that is less than 63 or external for an IS-IS metric that is greater than 64 and less than 128. The default is internal. The **match** keyword redistributes routes into other routing domains and is used with one of the following options: **external** [**1**|**2**] for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 1 or Type 2 external routes; **internal** for routes that are internal to a specific autonomous system; **nssa-external** [**1**|**2**] for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 1 or Type 2 external routes. The **tag** *tag-value* keyword-argument pair specifies the 32-bit decimal value that is 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 range from 0 to 4294967295. The **route-map** keyword specifies the route map to check for filtering the importing of routes from the source routing protocol to the current routing protocol. If this keyword is not specified, all routes are redistributed. If this keyword is specified, but no route map tags are listed, no routes are imported. The *map-tag* argument identifies a configured route map.

Removing the OSPF Configuration

To remove the entire OSPFv2 configuration that you have already enabled, enter the following commandperform the following steps:

Command	Purpose
clear configure router ospf *pid* **Example:** `hostname(config)# clear configure router ospf 1000`	Removes the entire OSPFv2 configuration that you have enabled. After the configuration is cleared, you must reconfigure OSPF using the **router ospf** command.

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, enter the following commandperform the following steps:

Command	Purpose
clear configure ipv6 router ospf *process-id* **Example:** `hostname(config)# clear configure ipv6 router ospf 1000`	Removes the entire OSPFv3 configuration that you have enabled. After the configuration is cleared, you must reconfigure OSPFv3 using the **ipv6 router ospf** command.

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 To enable OSPFv2, enter the following commands:

```
hostname(config)# router ospf 2
hostname(config-rtr)# network 10.0.0.0 255.0.0.0 area 0
```

Step 2 (Optional) To redistribute routes from one OSPFv2 process to another OSPFv2 process, enter the following commands:

```
hostname(config)# route-map 1-to-2 permit
hostname(config-route-map)# match metric 1
hostname(config-route-map)# set metric 5
hostname(config-route-map)# set metric-type type-1
hostname(config-route-map)# router ospf 2
hostname(config-rtr)# redistribute ospf 1 route-map 1-to-2
```

Step 3 (Optional) To configure OSPFv2 interface parameters, enter the following commands:

```
hostname(config)# router ospf 2
hostname(config-rtr)# network 10.0.0.0 255.0.0.0 area 0
hostname(config-rtr)# interface inside
hostname(config-interface)# ospf cost 20
hostname(config-interface)# ospf retransmit-interval 15
hostname(config-interface)# ospf transmit-delay 10
hostname(config-interface)# ospf priority 20
hostname(config-interface)# ospf hello-interval 10
hostname(config-interface)# ospf dead-interval 40
hostname(config-interface)# ospf authentication-key cisco
hostname(config-interface)# ospf message-digest-key 1 md5 cisco
hostname(config-interface)# ospf authentication message-digest
```

Step 4 (Optional) To configure OSPFv2 area parameters, enter the following commands:

```
hostname(config)# router ospf 2
hostname(config-rtr)# area 0 authentication
hostname(config-rtr)# area 0 authentication message-digest
hostname(config-rtr)# area 17 stub
hostname(config-rtr)# area 17 default-cost 20
```

Step 5 (Optional) To configure the route calculation timers and show the log neighbor up and down messages, enter the following commands:

```
hostname(config-rtr)# timers spf 10 120
hostname(config-rtr)# log-adj-changes [detail]
```

Step 6 (Optional) To show current OSPFv2 configuration settings, enter the **show ospf** command.

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

```
hostname(config)# show ospf

Routing Process "ospf 2" with ID 10.1.89.2 and Domain ID 0.0.0.2
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 5. Checksum Sum 0x 26da6
Number of opaque AS LSA 0. Checksum Sum 0x     0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
External flood list length 0
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        Area has no authentication
        SPF algorithm executed 2 times
        Area ranges are
        Number of LSA 5. Checksum Sum 0x 209a3
        Number of opaque link LSA 0. Checksum Sum 0x     0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Step 7 To clear the OSPFv2 configuration, enter the following command:

```
hostname(config)# clear configure router ospf pid
```

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 Examples 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.

The following example shows how to enable and configure OSPFv3 at the interface level:

```
hostname (config)# interface GigabitEthernet3/1
hostname (config-if)# ipv6 enable
hostname (config-if)# ipv6 ospf 1 area 1
```

The following is sample output from the **show running-config ipv6** command:

```
hostname (config)# show running-config ipv6
ipv6 router ospf 1
 log-adjacency-changes
```

The following is sample output from the **show running-config interface** command:

```
hostname (config-if)# show running-config interface GigabitEthernet3/1
interface GigabitEthernet3/1
 nameif fda
 security-level 100
 ip address 1.1.11.1 255.255.255.0 standby 1.1.11.2
 ipv6 address 9098::10/64 standby 9098::11
 ipv6 enable
 ipv6 ospf 1 area 1
```

The following examples show how to configure OSPFv3-specific interfaces:

```
hostname (config)# interface GigabitEthernet3/1
hostname (config-if)# nameif fda
hostname (config-if)# security-level 100
hostname (config-if)# ip address 10.1.11.1 255.255.255.0 standby 10.1.11.2
hostname (config-if)# ipv6 address 9098::10/64 standby 9098::11
hostname (config-if)# ipv6 enable
hostname (config-if)# ipv6 ospf cost 900
hostname (config-if)# ipv6 ospf hello-interval 20
hostname (config-if)# ipv6 ospf network broadcast
hostname (config-if)# ipv6 ospf database-filter all out
hostname (config-if)# ipv6 ospf flood-reduction
hostname (config-if)# ipv6 ospf mtu-ignore
hostname (config-if)# ipv6 ospf 1 area 1 instance 100
hostname (config-if)# ipv6 ospf encryption ipsec spi 890 esp null md5
12345678901234567890123456789012

hostname (config)# ipv6 router ospf 1
hostname (config)# area 1 nssa
hostname (config)# distance ospf intra-area 190 inter-area 100 external 100
hostname (config)# timers lsa arrival 900
hostname (config)# timers pacing flood 100
hostname (config)# timers throttle lsa 900 900 900
hostname (config)# passive-interface fda
hostname (config)# log-adjacency-changes
hostname (config)# redistribute connected metric 100 metric-type 1 tag 700
```

For an example of how to configure an OSPFv3 virtual link, see the following URL:

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080b8fd06.shtml

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, enter one of the following commands:

Command	Purpose
show ospf [*process-id* [*area-id*]]	Displays general information about OSPFv2 routing processes.
show ospf border-routers	Displays the internal OSPFv2 routing table entries to the ABR and ASBR.

Command	Purpose
show ospf [*process-id* [*area-id*]] **database**	Displays lists of information related to the OSPFv2 database for a specific router.
show ospf flood-list *if-name*	Displays a list of LSAs waiting to be flooded over an interface (to observe OSPF v2packet pacing). OSPFv2 update packets are automatically paced so they are not sent less than 33 milliseconds apart. Without pacing, some update packets could get lost in situations where the link is slow, a neighbor could not receive the updates quickly enough, or the router could run out of buffer space. For example, without pacing, packets might be dropped if either of the following topologies exist: • A fast router is connected to a slower router over a point-to-point link. • During flooding, several neighbors send updates to a single router at the same time. Pacing is also used between resends to increase efficiency and minimize lost retransmissions. You also can display the LSAs waiting to be sent out of an interface. Pacing enables OSPFv2 update and retransmission packets to be sent more efficiently. There are no configuration tasks for this feature; it occurs automatically.
show ospf interface [*if_name*]	Displays OSPFv2-related interface information.
show ospf neighbor [*interface-name*] [*neighbor-id*] [**detail**]	Displays OSPFv2 neighbor information on a per-interface basis.
show ospf request-list *neighbor if_name*	Displays a list of all LSAs requested by a router.
show ospf retransmission-list *neighbor if_name*	Displays a list of all LSAs waiting to be resent.
show ospf [*process-id*] **summary-address**	Displays a list of all summary address redistribution information configured under an OSPFv2 process.
show ospf [*process-id*] **traffic**	Displays a list of different types of packets being sent or received by a specific OSPFv2 instance.
show ospf [*process-id*] **virtual-links**	Displays OSPFv2-related virtual links information.
show route cluster	Displays additional OSPFv2 route synchronization information in clustering.

To monitor or display various OSPFv3 routing statistics, enter one of the following commands:

Command	Purpose
show ipv6 ospf [*process-id* [*area-id*]]	Displays general information about OSPFv3 routing processes.
show ipv6 ospf [*process-id*] **border-routers**	Displays the internal OSPFv3 routing table entries to the ABR and ASBR.
show ipv6 ospf [*process-id* [*area-id*]] **database** [**external** | **inter-area prefix** | **inter-area-router** | **network** | **nssa-external** | **router** | **area** | **as** | **ref-lsa** | [*destination-router-id*] [**prefix** *ipv6-prefix*] [*link-state-id*]] [**link** [**interface** *interface-name*] [**adv-router** *router-id*] | **self-originate**] [**internal**] [**database-summary**]	Displays lists of information related to the OSPFv3 database for a specific router.
show ipv6 ospf [*process-id* [*area-id*]] **events**	Displays OSPFv3 event information.
show ipv6 ospf [*process-id*] [*area-id*] **flood-list** *interface-type interface-number*	Displays a list of LSAs waiting to be flooded over an interface (to observe OSPFv3 packet pacing). OSPFv3 update packets are automatically paced so they are not sent less than 33 milliseconds apart. Without pacing, some update packets could get lost in situations where the link is slow, a neighbor could not receive the updates quickly enough, or the router could run out of buffer space. For example, without pacing, packets might be dropped if either of the following topologies exist: • A fast router is connected to a slower router over a point-to-point link. • During flooding, several neighbors send updates to a single router at the same time. Pacing is also used between retransmissions to increase efficiency and minimize lost retransmissions. You also can display the LSAs waiting to be sent out of an interface. Pacing enables OSPFv3 update and retransmission packets to be sent more efficiently. There are no configuration tasks for this feature; it occurs automatically.
show ipv6 ospf [*process-id*] [*area-id*] **interface** [*type number*] [**brief**]	Displays OSPFv3-related interface information.
show ipv6 ospf neighbor [*process-id*] [*area-id*] [*interface-type interface-number*] [*neighbor-id*] [**detail**]	Displays OSPFv3 neighbor information on a per-interface basis.
show ipv6 ospf [*process-id*] [*area-id*] **request-list** [*neighbor*] [*interface*] [*interface-neighbor*]	Displays a list of all LSAs requested by a router.

Command	Purpose
show ipv6 ospf [*process-id*] [*area-id*] **retransmission-list** [*neighbor*] [*interface*] [*interface-neighbor*]	Displays a list of all LSAs waiting to be resent.
show ipv6 ospf statistic [*process-id*] [**detail**]	Displays various OSPFv3 statistics.
show ipv6 ospf [*process-id*] **summary-prefix**	Displays a list of all summary address redistribution information configured under an OSPFv3 process.
show ipv6 ospf [*process-id*] **timers** [**lsa-group | rate-limit**]	Displays OSPFv3 timers information.
show ipv6 ospf [*process-id*] **traffic** [*interface_name*]	Displays OSPFv3 traffic-related statistics.
show ipv6 ospf virtual-links	Displays OSPFv3-related virtual links information.
show ipv6 route cluster [**failover**] [**cluster**] [**interface**] [**ospf**] [**summary**]	Displays the IPv6 routing table sequence number, IPv6 reconvergence timer status, and IPv6 routing entries sequence number in a cluster.

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 27-1 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 27-1 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 command: **route ospf** We introduced the following screen: Configuration > Device Setup > Routing > OSPF.

Table 27-1 Feature History for OSPF (continued)

Feature Name	Platform Releases	Feature Information
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. We introduced or modified the following commands: **show route cluster**, **show ipv6 route cluster**, **debug route cluster, router-id cluster-pool**.
OSPFv3 Support for IPv6		OSPFv3 routing is supported for IPv6. We introduced or modified the following commands: **ipv6 ospf**, **ipv6 ospf area**, **ipv6 ospf cost, ipv6 ospf database-filter all out, ipv6 ospf dead-interval, ipv6 ospf encryption, ipv6 ospf hello-interval, ipv6 ospf mtu-ignore, ipv6 ospf neighbor, ipv6 ospf network, ipv6 ospf flood-reduction, ipv6 ospf priority, ipv6 ospf retransmit-interval, ipv6 ospf transmit-delay, ipv6 router ospf, ipv6 router ospf area, ipv6 router ospf default, ipv6 router ospf default-information, ipv6 router ospf distance, ipv6 router ospf exit, ipv6 router ospf ignore, ipv6 router ospf log-adjacency-changes, ipv6 router ospf no, ipv6 router ospf passive-interface, ipv6 router ospf redistribute, ipv6 router ospf router-id, ipv6 router ospf summary-prefix, ipv6 router ospf timers, area encryption, area range, area stub, area nssa, area virtual-link, default, default-information originate, distance, ignore lsa mospf, log-adjacency-changes, redistribute, router-id, summary-prefix, timers lsa arrival, timers pacing flood, timers pacing lsa-group, timers pacing retransmission, timers throttle, show ipv6 ospf, show ipv6 ospf border-routers, show ipv6 ospf database, show ipv6 ospf events, show ipv6 ospf flood-list, show ipv6 ospf graceful-restart, show ipv6 ospf interface, show ipv6 ospf neighbor, show ipv6 ospf request-list, show ipv6 ospf retransmission-list, show ipv6 ospf statistic, show ipv6 ospf summary-prefix, show ipv6 ospf timers, show ipv6 ospf traffic, show ipv6 ospf virtual-links, show ospf, show running-config ipv6 router, clear ipv6 ospf, clear configure ipv6 router, debug ospfv3, ipv6 ospf neighbor**. 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 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 28-1
- Licensing Requirements for EIGRP, page 28-2
- Guidelines and Limitations, page 28-3
- Task List to Configure an EIGRP Process, page 28-3
- Configuring EIGRP, page 28-4
- Customizing EIGRP, page 28-7
- Monitoring EIGRP, page 28-28
- Configuration Example for EIGRP, page 28-29
- Feature History for EIGRP, page 28-30

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 24-9.

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.
- A maximum of one EIGRP process is 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 28-4 or the “Enabling EIGRP Stub Routing” section on page 28-6.

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 28-8 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 28-21 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 28-18 for more information.

Step 6 (Optional) Define static EIGRP neighbors on the Static Neighbor pane.

See the "Defining an EIGRP Neighbor" section on page 28-17 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 28-13 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 28-9 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 28-25 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 28-4
- Enabling EIGRP Stub Routing, page 28-6

Enabling EIGRP

You can only enable one EIGRP routing process on the ASA.

To enable EIGRP, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more **network** statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Configuring Interfaces for EIGRP" section on page 28-9.

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 28-4 and the "Enabling EIGRP Stub Routing" section on page 28-6 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 28-8 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:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more **network** statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the section "Configuring Passive Interfaces" section on page 28-12.

Step 3	**eigrp stub** {**receive-only** | [**connected**] [**redistributed**] [**static**] [**summary**]}
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
	`hostname(config-router)# eigrp stub {receive-only | [connected] [redistributed] [static] [summary]}`
		Configures the stub routing process. You must specify which networks are advertised by the stub routing process to the distribution router. Static and connected networks are not automatically redistributed into the stub routing process.

Note A stub routing process does not maintain a full topology table. At a minimum, stub routing needs a default route to a distribution router, which makes the routing decisions.

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 28-8
- Configuring Interfaces for EIGRP, page 28-9
- Configuring the Summary Aggregate Addresses on Interfaces, page 28-13
- Changing the Interface Delay Value, page 28-14
- Enabling EIGRP Authentication on an Interface, page 28-15

- Defining an EIGRP Neighbor, page 28-17
- Redistributing Routes Into EIGRP, page 28-18
- Filtering Networks in EIGRP, page 28-21
- Customizing the EIGRP Hello Interval and Hold Time, page 28-23
- Disabling Automatic Route Summarization, page 28-24
- Configuring Default Information in EIGRP, page 28-25
- Disabling EIGRP Split Horizon, page 28-26
- Restarting the EIGRP Process, page 28-27

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:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more **network** statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Configuring Passive Interfaces" section on page 28-12.

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 a **network** command that includes the network to which the interface is attached, and use the **passive-interface** command to prevent that interface from sending or receiving EIGRP updates.

To configure interfaces for EIGRP, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	`hostname(config-router)#` **network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Defining a Network for an EIGRP Routing Process" section on page 28-8.

Step 3	(Optional) Do one of the following to customize an interface to participate in EIGRP routing:	

Command	Purpose

no default-information {**in** | **out** | **WORD**}
 	
Example:
`hostname(config)# router eigrp 2`
`hostname(config-router)# network 10.0.0.0 255.0.0.0`
`hostname(config-router)# no default-information {in | out | WORD}`
	Allows you to control the sending or receiving of candidate default route information.
	
	Entering the **no default-information in** command causes the candidate default route bit to be blocked on received routes. Entering the **no default-information out** command disables the setting of the default route bit in advertised routes.
	
	See the "Configuring Default Information in EIGRP" section on page 28-25 for more information on this particular option.

authentication mode eigrp *as-num* **md5** **Example:** `hostname(config)# authentication mode eigrp 2 md5`	Enables MD5 authentication of EIGRP packets. The *as-num* argument is the autonomous system number of the EIGRP routing process configured on the ASA. If EIGRP is not enabled or if you enter the wrong number, the ASA returns the following error message: `% Asystem(100) specified does not exist` See the "Enabling EIGRP Authentication on an Interface" section on page 28-15 for more information on this particular option.
delay *value* **Example:** `hostname(config-if)# delay 200`	The *value* argument entered is in tens of microseconds. To set the delay for 2000 microseconds, you enter a *value* of 200. To view the delay value assigned to an interface, use the **show interface** command. See the "Changing the Interface Delay Value" section on page 28-14 for more information on this particular option.
hello-interval eigrp *as-num seconds* **Example:** `hostname(config)# hello-interval eigrp 2 60`	Allows you to change the hello interval. See the "Customizing the EIGRP Hello Interval and Hold Time" section on page 28-23 for more information on this particular option.
hold-time eigrp *as-num seconds* **Example:** `hostname(config)# hold-time eigrp 2 60`	Allows you to change the hold time. See the "Customizing the EIGRP Hello Interval and Hold Time" section on page 28-23 for more information on this particular option.

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:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	`hostname(config-router)#` **network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more **network** statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Defining a Network for an EIGRP Routing Process" section on page 28-8.

Step 3	**passive-interface** {**default** | *if-name*}
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
	`hostname(config-router)# passive-interface {default}`
		Prevents an interface from sending or receiving EIGRP routing message.
		
		Using the **default** keyword disables EIGRP routing updates on all interfaces. Specifying an interface name, as defined by the **nameif** command, disables EIGRP routing updates on the specified interface. You can use multiple **passive-interface** commands in your EIGRP router configuration.

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:

Detailed Steps

	Command	Purpose
Step 1	**interface** *phy_if* **Example:** `hostname(config)# interface inside`	Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.
Step 2	**summary-address eigrp** *as-num address mask* [*distance*] **Example:** `hostname(config-if)# summary-address eigrp 2 address mask [20]`	Creates the summary address. By default, EIGRP summary addresses that you define have an administrative distance of 5. You can change this value by specifying the optional *distance* argument in the **summary-address** command.

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 28-9.

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:

Detailed Steps

	Command	Purpose
Step 1	**interface** *phy_if* **Example:** `hostname(config)# interface inside`	Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.
Step 2	**delay** *value* **Example:** `hostname(config-if)# delay 200`	The *value* argument entered is in tens of microseconds. To set the delay for 2000 microseconds, you enter a *value* of 200. To view the delay value assigned to an interface, use the **show interface** command.

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 28-9.

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:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that falls within the defined network participate in the EIGRP routing process.
		
		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, see the "Configuring EIGRP" section on page 28-4.

Step 3	**interface** *phy_if* **Example:** `hostname(config)# interface inside`	Enters interface configuration mode for the interface on which you are configuring EIGRP message authentication.

	Command	Purpose
Step 4	**authentication mode eigrp** *as-num* **md5** **Example:** `hostname(config)# authentication mode eigrp 2 md5`	Enables MD5 authentication of EIGRP packets. The *as-num* argument is the autonomous system number of the EIGRP routing process configured on the ASA. If EIGRP is not enabled or if you enter the wrong number, the ASA returns the following error message: `% Asystem(100) specified does not exist`
Step 5	**authentication key eigrp** *as-num key* **key-id** *key-id* **Example:** `hostname(config)# authentication key eigrp 2 cisco key-id 200`	Configures the key used by the MD5 algorithm. The *as-num* argument is the autonomous system number of the EIGRP routing process configured on the ASA. If EIGRP is not enabled or if you enter the wrong number, the ASA returns the following error message: `% Asystem(100) specified does not exist%` The *key* argument can include up to 16 characters, including alphabets, numbers and special characters. **Note** White spaces are not allowed, in the key argurment. The *key-id* argument is a number that can range from 0 to 255.

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 28-9.

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:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**neighbor** *ip-addr* **interface** *if_name*
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# neighbor 10.0.0.0 interface interface1`
		Defines the static neighbor.
		
		The *ip-addr* argument is the IP address of the neighbor.
		
		The *if-name* argument is the name of the interface, as specified by the **nameif** command, through which that neighbor is available. You can define multiple neighbors for an EIGRP routing process.

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 26, "Defining Route Maps," for more information about creating a route map.

To redistribute routes into the EIGRP routing process, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	**default-metric** *bandwidth delay reliability loading mtu*
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# default-metric bandwidth delay reliability loading mtu`
		(Optional) Specifies the default metrics that should be applied to routes redistributed into the EIGRP routing process.
		
		If you do not specify a default metric in the EIGRP router configuration, you must specify the metric values in each **redistribute** command. If you specify the EIGRP metrics in the **redistribute** command and have the **default-metric** command in the EIGRP router configuration, the metrics in the **redistribute** command are used.

Step 3	Do one of the following to redistribute the selected route type into the EIGRP routing process:	

Command	Purpose
redistribute connected [**metric** *bandwidth delay reliability loading mtu*] [**route-map** *map_name*] **Example:** `hostname(config-router): redistribute connected [metric bandwidth delay reliability loading mtu] [route-map map_name]`	Redistributes connected routes into the EIGRP routing process. You must specify the EIGRP metric values in the **redistribute** command if you do not have a **default-metric** command in the EIGRP router configuration.
redistribute static [**metric** *bandwidth delay reliability loading mtu*] [**route-map** *map_name*] **Example:** `hostname(config-router): redistribute static [metric bandwidth delay reliability loading mtu] [route-map map_name]`	Redistributes static routes into the EIGRP routing process.
redistribute ospf *pid* [**match** {**internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**]}] [**metric** *bandwidth delay reliability loading mtu*] [**route-map** *map_name*] **Example:** `hostname(config-router): redistribute ospf pid [match {internal | external [1 | 2] | nssa-external [1 | 2]}] [metric bandwidth delay reliability loading mtu] [route-map map_name]`	Redistributes routes from an OSPF routing process into the EIGRP routing process.
redistribute rip [**metric** bandwidth delay reliability load mtu] [**route-map map_name**] **Example:** `(config-router): redistribute rip [metric bandwidth delay reliability load mtu] [route-map map_name]`	Redistributes routes from a RIP routing process into the EIGRP routing process.

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 26, “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 ACL that defines the routes that you want to advertise. That is, create a standard ACL that defines the routes that you want to filter from sending or receiving updates.

To filter networks in EIGRP, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	`hostname(config-router)#` **network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Configuring Interfaces for EIGRP" section on page 28-9.

Step 3	Do one of the following to filter networks sent or received in EIGRP routing updates:	

	distribute-list *acl* **out** [**connected** | **ospf** | **rip** | **static** | **interface** *if_name*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
	`hostname(config-router): distribute-list acl out [connected]`
		Filters networks sent in EIGRP routing updates.
		
		You can specify an interface to apply the filter to only those updates that are sent by that specific interface.
		
		You can enter multiple **distribute-list** commands in your EIGRP router configuration.

	distribute-list *acl* **in** [**interface** *if_name*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
	`hostname(config-router): distribute-list acl in [interface interface1]`
		Filters networks received in EIGRP routing updates.
		
		You can specify an interface to apply the filter to only those updates that are received by that 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 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 ACL 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:

Detailed Steps

	Command	Purpose
Step 1	**interface** *phy_if* **Example:** `hostname(config)# interface inside`	Enters interface configuration mode for the interface on which you are configuring the hello interval or advertised hold time.
Step 2	**hello-interval eigrp** *as-num seconds* **Example:** `hostname(config)# hello-interval eigrp 2 60`	Changes the hello interval.
Step 3	**hold-time eigrp** *as-num seconds* **Example:** `hostname(config)# hold-time eigrp 2 60`	Changes the hold time.

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, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.
Step 2	**no auto-summary** **Example:** `hostname(config-router)# no auto-summary`	You cannot configure this value. Automatic summary addresses have an administrative distance of 5.

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.

To configure default routing information, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router eigrp** *as-num* **Example:** `hostname(config)# router eigrp 2`	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The *as-num* argument is the autonomous system number of the EIGRP routing process.

Step 2	`hostname(config-router)#` **network** *ip-addr* [*mask*]
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
		Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
		
		Directly connected and static networks that fall within the defined network are advertised by the ASA. Additionally, only interfaces with an IP address that fall within the defined network participate in the EIGRP routing process.
		
		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, see the "Configuring Interfaces for EIGRP" section on page 28-9.

Step 3	**no default-information** {**in** | **out** | **WORD**}
	
	Example:
	`hostname(config)# router eigrp 2`
	`hostname(config-router)# network 10.0.0.0 255.0.0.0`
	`hostname(config-router)# no default-information {in | out | WORD}`
		Controls the sending or receiving of candidate default route information.
		
		Entering the **no default-information in** command causes the candidate default route bit to be blocked on received routes. Entering the **no default-information out** command 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:

Detailed Steps

	Command	Purpose
Step 1	**interface** *phy_if* **Example:** `hostname(config)# interface phy_if`	Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.
Step 2	**no split-horizon eigrp** *as-number* **Example:** `hostname(config-if)# no split-horizon eigrp 2`	Disables the split horizon.

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, enter the following commandperform the following steps:

Command	Purpose
clear eigrp pid {**1-65535** | **neighbors** | **topology** | **events**)} **Example:** `hostname(config)# clear eigrp pid 10 neighbors`	Restarts an EIGRP process or clears redistribution or counters.

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, enter one of the following commandsperform the following steps:

Command	Purpose
Monitoring EIGRP Routing	
router-id	Displays the router-id for this EIGRP process.
show eigrp [*as-number*] **events** [{*start end*} | **type**]	Displays the EIGRP event log.
show eigrp [*as-number*] **interfaces** [*if-name*] [**detail**]	Displays the interfaces participating in EIGRP routing.
show eigrp [*as-number*] **neighbors** [**detail** | **static**] [*if-name*]	Displays the EIGRP neighbor table.
show eigrp [*as-number*] **topology** [*ip-addr* [*mask*] | **active** | **all-links** | **pending** | **summary** | **zero-successors**]	Displays the EIGRP topology table.
show eigrp [*as-number*] **traffic**	Displays EIGRP traffic statistics.
show mfib cluster	Displays MFIB information in terms of forwarding entries and interfaces.
show route cluster	Displays additional route synchronization details for clustering.
Disabling EIGRP Logging Messages	
no eigrp log-neighbor-changes	Disables the logging of neighbor change messages. Enter this command in router configuration mode for the EIGRP routing process.
no eigrp log-neighbor-warnings	Disables the logging of neighbor warning messages.

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.

Configuration Example for EIGRP

The following example shows how to enable and configure EIGRP with various optional processes:

Step 1 To enable EIGRP, enter the following commands:

```
hostname(config)# router eigrp 2
hostname(config-router)# network 10.0.0.0 255.0.0.0
```

Step 2 To configure an interface from sending or receiving EIGRP routing messages, enter the following command:

```
hostname(config-router)# passive-interface {default}
```

Step 3 To define an EIGRP neighbor, enter the following command:

```
hostname(config-router)# neighbor 10.0.0.0 interface interface1
```

Step 4 To configure the interfaces and networks that participate in EIGRP routing, enter the following command:

```
hostname(config-router)# network 10.0.0.0 255.0.0.0
```

Step 5 To change the interface delay value used in EIGRP distance calculations, enter the following commands:

```
hostname(config-router)# exit
hostname(config)# interface phy_if
hostname(config-if)# delay 200
```

Feature History for EIGRP

Table 28-1 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 28-1 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 command: **route 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. We introduced or modified the following commands: **show route cluster**, **debug route cluster, show mfib cluster, debug mfib cluster**.



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 29-1
- Licensing Requirements for RIP, page 29-3
- Guidelines and Limitations, page 29-3
- Configuring RIP, page 29-4
- Customizing RIP, page 29-5
- Monitoring RIP, page 29-16
- Configuration Example for RIP, page 29-17
- Feature History for RIP, page 29-18

Information About RIP

This section includes the following topics:

- Routing Update Process, page 29-2
- RIP Routing Metric, page 29-2
- RIP Stability Features, page 29-2
- RIP Timers, page 29-2
- Using Clustering, page 29-3

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 24-9.

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 29-5 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 25-7 and then define a route map. For information, see the "Defining a Route Map" section on page 26-4.

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.

To enable the RIP routing process, enter the following command:

Command	Purpose

`router rip`
 	
Example:
`hostname(config)# router rip`
	Starts the RIP routing process and places you in router configuration mode.
	
	Use the **no router rip** command to remove the entire RIP configuration that you have enabled. After the configuration is cleared, you must reconfigure RIP using the **router rip** command.

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 29-10.
- Enable RIP version. See the "Configuring the RIP Version" section on page 29-5.
- Enable default information origination.
- Define an IP Address for a Network to Add. See the "Filtering Networks in RIP" section on page 29-11.
- Configure an Interface. See the "Configuring Interfaces for RIP" section on page 29-7.

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 29-4.

Customizing RIP

This section describes how to configure RIP and includes the following topics:

- Configuring the RIP Version, page 29-5
- Configuring Interfaces for RIP, page 29-7
- Configuring the RIP Send and Receive Version on an Interface, page 29-9
- Configuring Route Summarization, page 29-10
- Filtering Networks in RIP, page 29-11
- Redistributing Routes into the RIP Routing Process, page 29-13
- Enabling RIP Authentication, page 29-15
- Restarting the RIP Process, page 29-16

Configuring the RIP Version

To specify the version of RIP used by the ASA, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router rip** **Example:** `hostname(config)# router rip`	Starts the RIP routing process and places you in router configuration mode.

Step 2	**network** *network_address*
	
	Example:
	`hostname(config)# router rip`
	`hostname(config-router)# network 10.0.0.0`
		Specifies the interfaces that will participate in the RIP routing process.
		
		If an interface belongs to a network defined by this command, the interface will participate in the RIP routing process. If an interface does not belong to a network defined by this command, the interface will not send or receive RIP updates.

Command	Purpose
Step 3 Enter one of the following numbers to customize an interface to participate in RIP routing:	

`version [1 | 2]`
 	
Example:
`hostname(config-router):# version [1]`
	Specifies the version of RIP used by the ASA.
	You can override this setting on a per-interface basis.
	In this example, Version 1 is entered.

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 29-7. 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 (using the **network** command) that includes the network to which the interface is attached, and configure the passive interfaces (using the **passive-interface** command) to prevent that interface from using RIP. Additionally, you can specify the version of RIP that is used by the ASA for updates.

To configure interfaces for RIP, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router rip** **Example:** `hostname(config)# router rip`	Starts the RIP routing process and places you in router configuration mode.

Step 2	**network** *network_address*
	
	Example:
	`hostname(config)# router rip`
	`hostname(config-router)# network 10.0.0.0`
		Specifies the interfaces that will participate in the RIP routing process.
		
		If an interface belongs to a network defined by this command, the interface will participate in the RIP routing process. If an interface does not belong to a network defined by this command, it will not send or receive RIP updates.

Step 3	**passive-interface** [**default** | *if_name*] **Example:** `hostname(config-router)# passive-interface [default]`	Specifies an interface to operate in passive mode. Using the **default** keyword causes all interfaces to operate in passive mode. Specifying an interface name sets only that interface to passive mode. In passive mode, RIP routing updates are accepted by, but not sent out of, the specified interface. You can enter this command for each interface that you want to set to passive mode.

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 29-8.

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:

Detailed Steps

	Command	Purpose
Step 1	**interface** *phy_if* **Example:** `hostname(config)# interface phy_if`	Enters interface configuration mode for the interface that you are configuring.
Step 2	Do one of the following to send or receive RIP updates on a per-interface basis.	
	rip send version {[**1**] [**2**]} **Example:** `hostname(config-if)# rip send version 1`	Specifies the version of RIP to use when sending RIP updates out of the interface. In this example, Version 1 is selected.
	rip receive version {[**1**] [**2**]} **Example:** `hostname(config-if)# rip receive version 2`	Specifies the version of RIP advertisements permitted to be received by an interface. In this example, Version 2 is selected. RIP updates received on the interface that do not match the allowed version are dropped.

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.

To disable automatic route summarization, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**router rip** **Example:** `hostname(config)# router rip`	Enables the RIP routing process and places you in router configuration mode.
Step 2	**no auto-summarize** **Example:** `hostname(config-router):# no auto-summarize`	Disables automatic route summarization.

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 ACL 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.

Detailed Steps

	Command	Purpose
Step 1	**router rip** **Example:** `hostname(config)# router rip`	Enables the RIP routing process and places you in router configuration mode.

Step 2	**distribute-list** *acl* **in** [**interface** *if_name*]
	distribute-list *acl* **out** [**connected** | **eigrp** | **interface** *if_name* | **ospf** | **rip** | **static**]
	
	Example:
	`hostname(config-router)# distribute-list acl2 in [interface interface1]`
	`hostname(config-router)# distribute-list acl3 out [connected]`
		Filters the networks sent in updates.
		
		You can specify an interface to apply the filter to only those updates that are received or sent by that interface. You can enter this command for each interface to which you want to apply a filter. If you do not specify an interface name, the filter is applied to all RIP updates.

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 29-12.

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 29-12.

Adding or Editing a Filter Rule

After you have configured a filter rule (see the "Filtering Networks in RIP" section on page 29-11), 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 26, "Defining a Route Map," for more information about creating a route map.

To redistribute a route into the RIP routing process, enter one of the following commands:

Command	Purpose
redistribute connected [**metric** *metric-value* | **transparent**] [**route-map** *route-map-name*] **Example:** `hostname(config-router): # redistribute connected [metric metric-value | transparent] [route-map route-map-name]`	Redistributes connected routes into the RIP routing process. You must specify the RIP metric values in the **redistribute** command if you do not have a **default-metric** command in the RIP router configuration.
redistribute static [**metric** {*metric_value* | **transparent**}] [**route-map** *map_name*] **Example:** `hostname(config-router):# redistribute static [metric {metric_value | transparent}] [route-map map_name]`	Redistributes static routes into the EIGRP routing process.
redistribute ospf *pid* [**match** {**internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**]}] [**metric** {*metric_value* | **transparent**}] [**route-map** *map_name*] **Example:** `hostname(config-router):# redistribute ospf pid [match {internal | external [1 | 2] | nssa-external [1 | 2]}] [metric {metric_value | transparent}] [route-map map_name]`	Redistributes routes from an OSPF routing process into the RIP routing process.
redistribute eigrp *as-num* [**metric** {*metric_value* | **transparent**}] [**route-map** *map_name*] **Example:** `hostname(config-router):# redistribute eigrp as-num [metric {metric_value | transparent}] [route-map map_name]`	Redistributes routes from an EIGRP routing process into the RIP routing process.

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 26-4.

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:

Detailed Steps

	Command	Purpose
Step 1	**router rip** *as-num* **Example:** `hostname(config)# router rip 2`	Creates the RIP routing process and enters router configuration mode for this RIP process. The *as-num* argument is the autonomous system number of the RIP routing process.
Step 2	**interface** *phy_if* **Example:** `hostname(config)# interface phy_if`	Enters interface configuration mode for the interface on which you are configuring RIP message authentication.
Step 3	**rip authentication mode** {**text** | **md5**} **Example:** `hostname(config-if)# rip authentication mode md5`	Sets the authentication mode. By default, text authentication is used. We recommend that you use MD5 authentication.
Step 4	**rip authentication key** *key* **key-id** *key-id* **Example:** `hostname(config-if)# rip authentication key cisco key-id 200`	Configures the authentication key used by the MD5 algorithm. The *key* argument can include up to 16 characters. The *key-id* argument is a number from 0 to 255.

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 one of the following options:

- **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, enter the following commandperform the following steps:

Command	Purpose
clear rip *pid* {**process** | **redistribution** | **counters** [**neighbor** [*neighbor-interface*] [*neighbor-id*]]}	Removes the entire RIP configuration that you have enabled. After the configuration is cleared, you must reconfigure RIP again using the **router rip** command.
Example: `hostname(config)# clear rip`	

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > RIP > Setup**.

Step 2 Click **Reset**.

Monitoring RIP

We recommend that you only use the **debug** commands to troubleshoot specific problems or during troubleshooting sessions with the Cisco TAC.

Debugging output is assigned high priority in the CPU process and can render the ASA unusable. It is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect performance. For examples and descriptions of the command output, see the command reference.

To monitor or debug various RIP routing statistics, enter one of the following commands:

Command	Purpose
Monitoring RIP Routing	
show rip database	Display the contents of the RIP routing database.
show running-config router rip	Displays the RIP commands.
show route cluster	Displays additional route synchronization details for clustering.
Debugging RIP	
debug rip events	Displays RIP processing events.
debug rip database	Displays RIP database events.
debug route cluster	Enables RIB table replication trace messages to determine if the RIB is correctly synchronized to the slave units in clustering.

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:

```
hostname(config)# router rip 2
hostname(config-router)# default-information originate
hostname(config-router)# version [1]
hostname(config-router)# network 225.25.25.225
hostname(config-router)# passive-interface [default]
hostname(config-router)# redistribute connected [metric bandwidth delay reliability
loading mtu] [route-map map_name]
```

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 26-4.

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 26-4.

Step 14 Click **OK**.

Feature History for RIP

Table 29-1 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 29-1 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 **route rip** command. 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. We introduced or modified the following commands: **show route cluster**, **debug route cluster, show mfib cluster, debug mfib cluster**.



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 30-1
- Licensing Requirements for Multicast Routing, page 30-3
- Guidelines and Limitations, page 30-3
- Enabling Multicast Routing, page 30-4
- Customizing Multicast Routing, page 30-4
- Configuration Example for Multicast Routing, page 30-23
- Additional References, page 30-25
- Feature History for Multicast Routing, page 30-25

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 30-2
- PIM Multicast Routing, page 30-2

- Multicast Group Concept, page 30-2
- Clustering, page 30-3

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 Rendezvous Point, use the untranslated outside address of the ASA as the Rendezvous Point 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 30-19.

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, enter the following commandperform the following steps:

Command	Purpose

`multicast-routing`
 	
Example:
`hostname(config)# multicast-routing`
	Enables multicast routing.
	
	The number of entries in the multicast routing tables are limited by the amount of RAM on the ASA.

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 30-1 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 30-1 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 30-5
- Configuring a Static Multicast Route, page 30-6
- Configuring IGMP Features, page 30-7
- Configuring PIM Features, page 30-14
- Configuring a Multicast Group, page 30-19

- Configuring a Bidirectional Neighbor Filter, page 30-21
- Configuring a Multicast Boundary, page 30-22

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, enter the following command from the interface attached to the stub areaperform the following steps:

Command	Purpose
igmp forward interface *if_name* **Example:** `hostname(config-if)# igmp forward interface interface1`	Configures stub multicast routing and forwards IGMP messages.

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, enter one of the following commandsperform the following steps:

Command	Purpose
mroute *src_ip src_mask* {*input_if_name* | *rpf_neighbor*} [*distance*] **Example:** hostname(config)# mroute *src_ip src_mask* {*input_if_name* | *rpf_neighbor*} [*distance*]	Configures a static multicast route.
mroute *src_ip src_mask input_if_name* [**dense** *output_if_name*] [*distance*] **Example:** hostname(config)# mroute *src_ip src_mask input_if_name* [dense *output_if_name*] [*distance*]	Configures a static multicast route for a stub area. The **dense** *output_if_name* keyword and argument pair is only supported for stub multicast routing.

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 30-8
- Configuring IGMP Group Membership, page 30-8
- Configuring a Statically Joined IGMP Group, page 30-9
- Controlling Access to Multicast Groups, page 30-10
- Limiting the Number of IGMP States on an Interface, page 30-11
- Modifying the Query Messages to Multicast Groups, page 30-12
- Changing the IGMP Version, page 30-13

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, enter the following commandperform the following steps:

Command	Purpose

`no igmp`
 	
Example:
`hostname(config-if)# no igmp`
	Disables IGMP on an interface.
	
	To reenable IGMP on an interface, use the **igmp** command.

Note Only the **no igmp** command appears in the interface configuration.

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 30-9.

To have the ASA join a multicast group, enter the following commandperform the following steps:

Command	Purpose
igmp join-group *group-address* **Example:** `hostname(config-if)# igmp join-group mcast-group`	Configures the ASA to be a member of a multicast group. The *group-address* argument is the IP address of the group.

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.

Enter the **igmp static-group** command. The ASA does not accept the multicast packets, but instead forwards them to the specified interface.

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, enter the following commandperform the following steps:

Command	Purpose
igmp static-group **Example:** `hostname(config-if)# igmp static-group group-address`	Configures the ASA statically to join a multicast group on an interface. The *group-address* argument is the IP address of the group.

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:

Detailed Steps

	Command	Purpose
Step 1	`Do one of the following to create a standard or extended ACL:`	
	access-list *name* **standard** [**permit** | **deny**] *ip_addr mask* **Example:** `hostname(config)# access-list acl1 standard permit 192.52.662.25`	Creates a standard ACL for the multicast traffic. You can create more than one entry for a single ACL. You can use extended or standard ACLs. The *ip_addr mask* argument is the IP address of the multicast group being permitted or denied.
	access-list *name* **extended** [**permit** | **deny**] *protocol src_ip_addr src_mask dst_ip_addr dst_mask* **Example:** `hostname(config)# access-list acl2 extended permit protocol src_ip_addr src_mask dst_ip_addr dst_mask`	Creates an extended ACL. The *dst_ip_addr* argument is the IP address of the multicast group being permitted or denied.
Step 2	**igmp access-group** *acl* **Example:** `hostname(config-if)# igmp access-group acl`	Applies the ACL to an interface. The *acl* argument is the name of a standard or extended IP ACL.

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, enter the following commandperform the following steps:

Command	Purpose
igmp limit *number* **Example:** `hostname(config-if)# igmp limit 50`	Limits the number of IGMP states on an interface. Valid values range from 0 to 500, with 500 being the default value. Setting this value to 0 prevents learned groups from being added, but manually defined memberships (using the **igmp join-group** and **igmp static-group** commands) are still permitted. The **no** form of this command restores the default value.

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

Note The **igmp query-timeout** and **igmp query-interval** commands require IGMP Version 2.

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:

Detailed Steps

	Command	Purpose
Step 1	**igmp query-interval** *seconds* **Example:** `hostname(config-if)# igmp query-interval 30`	Sets the query interval time in seconds. Valid values range from 0 to 500; 125 is the default value. If the ASA does not hear a query message on an interface for the specified timeout value (by default, 255 seconds), then the ASA becomes the designated router and starts sending the query messages.
Step 2	**igmp query-timeout** *seconds* **Example:** `hostname(config-if)# igmp query-timeout 30`	Changes the timeout value of the query. `Valid values range from 0 to 500; 225 is the default value.`
Step 3	**igmp query-max-response-time** *seconds* **Example:** `hostname(config-if)# igmp query-max-response-time 30`	Changes the maximum query response time.

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 such as the **igmp query-timeout** and **igmp query-interval** commands.

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, enter the following commandperform the following steps:

Command	Purpose
igmp version {**1** | **2**} **Example:** `hostname(config-if)# igmp version 2`	Controls the version of IGMP that you want to run on the interface.

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 30-14
- Configuring a Static Rendezvous Point Address, page 30-15
- Configuring the Designated Router Priority, page 30-16
- Configuring and Filtering PIM Register Messages, page 30-17
- Configuring PIM Message Intervals, page 30-18
- Configuring a Route Tree, page 30-18
- Filtering PIM Neighbors, page 30-19

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:

Detailed Steps

	Command	Purpose
Step 1	**pim** **Example:** `hostname(config-if)# pim`	Enables or reenables PIM on a specific interface.
Step 2	**no pim** **Example:** `hostname(config-if)# no pim`	Disables PIM on a specific interface.

Note Only the **no pim** command appears in the interface configuration.

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 must use the **pim rp-address** command to specify the RP address.

You can configure the ASA to serve as RP to more than one group. The group range specified in the ACL determines the PIM RP group mapping. If an ACL 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, enter the following commandperform the following steps:

Command	Purpose
pim rp-address *ip_address* [*acl*] [**bidir**] **Example:** `hostname(config)# pim rp-address 10.86.75.23 [acl1] [bidir]`	Enables or reenables PIM on a specific interface. The *ip_address* argument is the unicast IP address of the router assigned to be a PIM RP. The *acl* argument is the name or number of a standard ACL that defines with which multicast groups the RP should be used. Do not use a host ACL with this command. Excluding the **bidir** keyword causes the groups to operate in PIM sparse mode.

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 30-19.

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, enter the following commandperform the following steps:

Command	Purpose
pim dr-priority *num* **Example:** `hostname(config-if)# pim dr-priority 500`	Changes the designated router priority. The *num* argument can be any number ranging from 1 to 4294967294.

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, enter the following commandperform the following steps:

Command	Purpose
pim accept-register {**list** *acl* | **route-map** *map-name*} **Example:** `hostname(config)# pim accept-register {list acl1 | route-map map2}`	Configures the ASA to filter PIM register messages. In the example, the ASA filters PIM register messages *acl1* and route map *map2*.

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:

Detailed Steps

	Command	Purpose
Step 1	**pim hello-interval** *seconds* **Example:** `hostname(config-if)# pim hello-interval 60`	Sends router query messages. Valid values for the *seconds* argument range from 1 to 3600 seconds.
Step 2	**pim join-prune-interval** *seconds* **Example:** `hostname(config-if)# pim join-prune-interval 60`	Changes the amount of time (in seconds) that the ASA sends PIM join or prune messages. Valid values for the *seconds* argument range from 10 to 600 seconds.

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 30-19.

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:

Detailed Steps

	Command	Purpose
Step 1	**access-list pim_nbr deny** *router-IP_addr PIM neighbor* **Example:** `hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255`	Uses a standard ACL to define the routers that you want to have participate in PIM. In the example, the following ACL, when used with the **pim neighbor-filter** command, prevents the 10.1.1.1 router from becoming a PIM neighbor.

Step 2	**pim neighbor-filter pim_nbr**
	
	Example:
	`hostname(config)# interface GigabitEthernet0/3`
	`hostname(config-if)# pim neighbor-filter pim_nbr`
		Filters neighbor routers.
		
		In the example, the 10.1.1.1 router is prevented from becoming a PIM neighbor on interface GigabitEthernet0/3.

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:

Detailed Steps

	Command	Purpose
Step 1	**access-list pim_nbr deny** *router-IP_addr PIM neighbor* **Example:** `hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255`	Uses a standard ACL to define the routers that you want to have participate in PIM. In the example, the following ACL, when used with the **pim neighbor-filter** command, prevents the 10.1.1.1 router from becoming a PIM neighbor.

Step 2	**pim bidirectional-neighbor-filter pim_nbr**
	
	Example:
	`hostname(config)# interface GigabitEthernet0/3`
	`hostname(config-if)# pim bidirectional neighbor-filter pim_nbr`
		Filters neighbor routers.
		
		In the example, the 10.1.1.1 router is prevented from becoming a PIM bidirectional neighbor on interface GigabitEthernet0/3.

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 entering the **multicast boundary** command. 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 by entering the **filter-autorp** keyword. 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, enter the following commandperform the following steps:

Command	Purpose
multicast boundary *acl* [**filter-autorp**] **Example:** `hostname(config-if)# multicast boundary acl1 [filter-autorp]`	Configures a multicast boundary.

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 Enable multicast routing:

```
hostname(config)# multicast-routing
```

Step 2 Configure a static multicast route:

```
hostname(config)# mroute src_ip src_mask {input_if_name | rpf_neighbor} [distance]
hostname(config)# exit
```

Step 3 Configure the ASA to be a member of a multicast group:

```
hostname(config)# interface
hostname(config-if)# igmp join-group group-address
```

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 30-25
- RFCs, page 30-25

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 30-2 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 30-2 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 **multicast-routing** command. We introduced the following screen: Configuration > Device Setup > Routing > Multicast.
Clustering support	9.0(1)	Support was added for clustering. We introduced the following commands: **debug mfib cluster**, **show mfib cluster**.



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 31-1
- Licensing Requirements for IPv6 Neighbor Discovery, page 31-5
- Prerequisites for IPv6 Neighbor Discovery, page 31-5
- Guidelines and Limitations, page 31-5
- Default Settings for IPv6 Neighbor Discovery, page 31-7
- Configuring IPv6 Neighbor Discovery, page 31-7
- Viewing and Clearing Dynamically Discovered Neighbors, page 31-18
- Monitoring IPv6 Neighbor Discovery, page 31-18
- Additional References, page 31-18
- Feature History for IPv6 Neighbor Discovery, page 31-20

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 31-2
- Neighbor Reachable Time, page 31-3
- Duplicate Address Detection, page 31-3
- Router Advertisement Messages, page 31-3
- Static IPv6 Neighbors, page 31-5

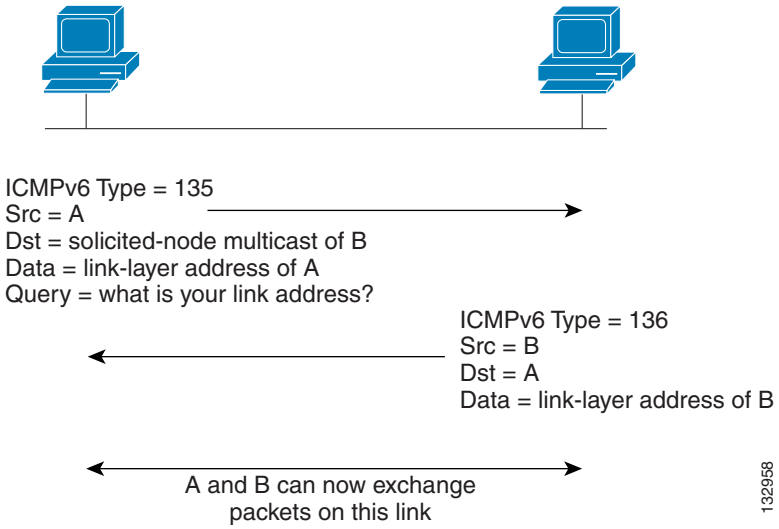
Neighbor Solicitation Messages

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 31-1 shows the neighbor solicitation and response process.

Figure 31-1 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.

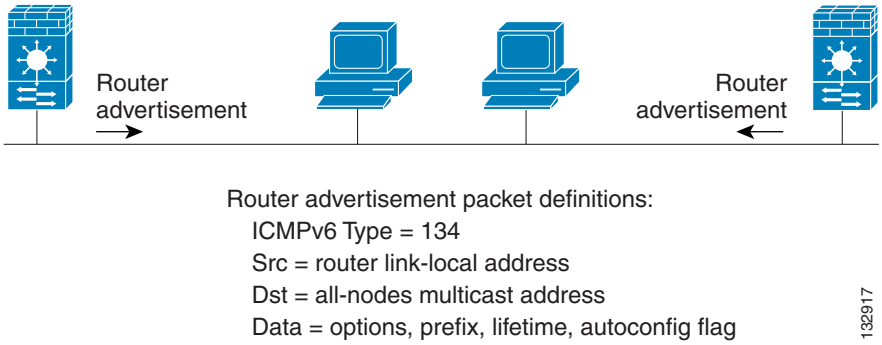
If the link-local address for an interface changes, Duplicate Address Detection is performed on the new link-local address and all of the other IPv6 address associated with the interface are regenerated (Duplicate Address Detection is performed only on the new link-local address).

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 31-2 shows an example of how to send router advertisement messages on an IPv6 configured interface.

Figure 31-2 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 13-15.

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 31-1 lists the default settings for IPv6 Neighbor Discovery.

Table 31-1 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

- Entering Interface Configuration Mode, page 31-7
- Configuring the Neighbor Solicitation Message Interval, page 31-8
- Configuring the Neighbor Reachable Time, page 31-9
- Configuring the Router Advertisement Transmission Interval, page 31-10
- Configuring the Router Lifetime Value, page 31-11
- Configuring DAD Settings, page 31-12
- Suppressing Router Advertisement Messages, page 31-13
- Configuring Address Config Flags for IPv6 DHCP Relay, page 31-14
- Configuring the IPv6 Prefix in Router Advertisements, page 31-15
- Configuring a Static IPv6 Neighbor, page 31-17

Entering Interface Configuration Mode

Configure neighbor discovery settings per interface. To enter interface configuration mode, perform the following steps.

Detailed Steps

Command	Purpose

interface *name*
 	
Example:
`hostname(config)# interface gigabitethernet 0/0`
`hostname(config-if)#`
	Enters interface configuration mode.

Configuring the Neighbor Solicitation Message Interval

To configure the interval between IPv6 neighbor solicitation retransmissions on an interface, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd ns-interval *value* **Example:** `hostname (config-if)# ipv6 nd ns-interval 9000`	Sets the interval between IPv6 neighbor solicitation retransmissions on an interface. Valid values for the value argument range from 1000 to 3600000 milliseconds. This information is also sent in router advertisement messages.

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 13-15 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.

Examples

The following example configures an IPv6 neighbor solicitation transmission interval of 9000 milliseconds for GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd ns-interval 9000
```

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, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd reachable-time *value* **Example:** `hostname (config-if)# ipv6 nd reachable-time 1700000`	Sets the amount of time that a remote IPv6 node is reachable. Valid values for the *value* argument range from 0 to 3600000 milliseconds. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

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 13-15.

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.

Examples

The following example configures an IPv6 reachable time of 1700000 milliseconds for the selected interface, GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd reachable-time 1700000
```

Configuring the Router Advertisement Transmission Interval

To configure the interval between IPv6 router advertisement transmissions on an interface, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd ra-interval [**msec**] *value* **Example:** `hostname (config-if)# ipv6 nd ra-interval 201`	Sets the interval between IPv6 router advertisement transmissions. The optional **msec** keyword indicates that the value provided is in milliseconds. If this keyword is not present, the value provided is in seconds. Valid values for the *value* argument range from 3 to 1800 seconds or from 500 to 1800000 milliseconds if the **msec** keyword is provided. 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. For more information, see the "Configuring the Router Lifetime Value" section on page 31-11. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value.

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 13-15.

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.

Examples

The following example configures an IPv6 router advertisement interval of 201 seconds for the selected interface, GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd ra-interval 201
```

Configuring the Router Lifetime Value

To configure the router lifetime value in IPv6 router advertisements on an interface, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd ra-lifetime [**msec**] *value* **Example:** `hostname (config-if)# ipv6 nd ra-lifetime 2000`	Specifies the length of time that nodes on the local link should consider the ASA as the default router on the link. The optional **msec** keyword indicates that the value provided is in milliseconds. If this keyword is not present, the value provided is in seconds. Valid values for the *value* argument range from 0 to 9000 seconds. Entering 0 indicates that the ASA should not be considered a default router on the selected interface.

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 13-15.

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.

Examples

The following example configures an IPv6 router lifetime value of 2000 seconds for the selected interface, GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd ra-lifetime 2000
```

Configuring DAD Settings

To specify DAD settings on the interface, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd dad attempts *value* **Example:** `hostname (config-if)# ipv6 nd dad attempts 20`	Specifies 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. Valid values for the *value* argument range from 0 to 600. A zero value disables DAD processing on the specified interface.

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 13-15.

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.

Examples

The following example configures a DAD attempt value of 20 for the selected interface, GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd dad attempts 20
```

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, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd suppress-ra *seconds* **Example:** `hostname (config-if)# ipv6 nd suppress-ra 900`	Suppresses the router lifetime value. The *seconds* argument specifies the validity of the ASA as a default router on this interface. Valid values range from 0 to 9000 seconds. A zero indicates that the ASA should not be considered a default router on the specified interface. Entering this command causes the ASA to appear as a regular IPv6 neighbor on the link and not as an IPv6 router.

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.

Examples

The following example suppresses an IPv6 router advertisement transmission for the specified interface, which is GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd suppress-ra 900
```

Configuring Address Config Flags for IPv6 DHCP Relay

You can add a flag to IPv6 router advertisements to inform IPv6 autoconfiguration clients to use DHCPv6 to obtain an IPv6 address and/or additional information such as the DNS server address.

Detailed Steps

Command	Purpose
ipv6 nd managed-config-flag **Example:** `hostname (config-if)# ipv6 nd managed-config-flag`	Sets 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.
ipv6 nd other-config-flag **Example:** `hostname (config-if)# ipv6 nd other-config-flag`	Sets 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.

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, enter the following commandperform the following steps.

Detailed Steps

Command	Purpose
ipv6 nd prefix *ipv6-prefix/prefix-length* | **default** [[*valid-lifetime preferred-lifetime*] | [**at** *valid-date preferred-date*] | **infinite** | **no-advertise** | **off-link** | **no-autoconfig**] **Example:** `hostname (config-if)# ipv6 nd prefix 2001:DB8::/32 1000 900`	Configures which IPv6 prefixes are included in IPv6 router advertisements. The prefix advertisement can be used by neighboring devices to autoconfigure their interface addresses. Stateless autoconfiguration uses IPv6 prefixes provided in router advertisement messages to create the global unicast address from the link-local address. The **at** *valid-date preferred-date* syntax indicates the date and time at which the lifetime and preference expire. The prefix is valid until this specified date and time are reached. Dates are expressed in the form *date-valid-expire month-valid-expire hh:mm-valid-expire date-prefer-expire month-prefer-expire hh:mm-prefer-expire*. The **default** keyword indicates that default values are used. The optional **infinite** keyword specifies that the valid lifetime does not expire. The *ipv6-prefix* argument specifies the IPv6 network number to include in router advertisements. This argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons. The optional **no-advertise** keyword indicates to hosts on the local link that the specified prefix is not to be used for IPv6 autoconfiguration. The optional **no-autoconfig** keyword indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration. The optional **off-link** keyword indicates that the specified prefix is not used for on-link determination. The *preferred-lifetime* argument specifies the amount of time (in seconds) that the specified IPv6 prefix is advertised as being preferred. Valid values range from 0 to 4294967295 seconds. The maximum value represents infinity, which can also be specified with infinite. The default is 604800 (7 days). The *prefix-length* argument specifies the length of the IPv6 prefix. This value indicates how many of the high-order, contiguous bits of the address comprise the network portion of the prefix. The slash (/) must precede the prefix length. The *valid-lifetime* argument specifies the amount of time that the specified IPv6 prefix is advertised as being valid. Valid values range from 0 to 4294967295 seconds. The maximum value represents infinity, which can also be specified with infinite. The default is 2592000 (30 days).

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.

Examples

The following example includes the IPv6 prefix 2001:DB8::/32, with a valid lifetime of 1000 seconds and a preferred lifetime of 900 seconds, in router advertisements sent out on the specified interface, which is GigabitEthernet 0/0:

```
hostname (config)# interface gigabitethernet 0/0
hostname (config-if)# ipv6 nd prefix 2001:DB8::/32 1000 900
```

Configuring a Static IPv6 Neighbor

To configure a static entry in the IPv6 neighbor discovery cache, enter the following command.

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 13-15.

To add an IPv6 static neighbor, perform the following steps.

Detailed Steps

Command	Purpose
ipv6 neighbor *ipv6_address if_name mac_address* **Example:** `hostname(config-if)# ipv6 neighbor 3001:1::45A inside 002.7D1A.9472`	Configures a static entry in the IPv6 neighbor discovery cache. The *ipv6_address* argument is the link-local IPv6 address of the neighbor, the *if_name* argument is the interface through which the neighbor is available, and the *mac_address* argument is the MAC address of the neighbor interface.

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.

Examples

The following example adds a static entry for an inside host with an IPv6 address of 3001:1::45A and a MAC address of 002.7D1a.9472 to the neighbor discovery cache:

```
hostname(config-if)# ipv6 neighbor 3001:1::45A inside 002.7D1A.9472
```

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.

Monitoring IPv6 Neighbor Discovery

To monitor IPv6 neighbor discovery parameters, enter the following command:

Command	Purpose
`show ipv6 interface`	Displays the usability status of interfaces configured for IPv6. Including the interface name, such as “outside” and displays the settings for the specified interface. Excludes the name from the command and displays the settings for all interfaces that have IPv6 enabled on them. Output for the command shows the following: • The name and status of the interface. • The link-local and global unicast addresses. • The multicast groups to which the interface belongs. • ICMP redirect and error message settings. • Neighbor discovery settings. • The actual time when the command is set to 0. • The neighbor discovery reachable time that is being used.

Additional References

For additional information related to implementing IPv6 prefixes, see the following topics:

- Related Documents for IPv6 Prefixes, page 31-20

- RFCs for IPv6 Prefixes and Documentation, page 31-20

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 31-2 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 31-2 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. We introduced the following commands: **ipv6 nd ns-interval**, **ipv6 nd ra-lifetime**, **ipv6 nd suppress-ra**, **ipv6 neighbor**, **ipv6 nd prefix**, **ipv6 nd dad-attempts**, **ipv6 nd reachable-time**, **ipv6 address**, **ipv6 enforce-eui64**.
Address Config Flags for IPv6 DHCP Relay	9.0(1)	We modified the following screen: Configuration > Device Device Setup > Interfaces > IPv6. We introduced the following commands: **ipv6 nd managed-config-flag**, **ipv6 nd other-config-flag**.



PART 7

Configuring AAA Servers and the Local Database



Information About AAA

This chapter describes authentication, authorization, and accounting (AAA, pronounced “triple A”). AAA is a a set of services for controlling access to computer resources, enforcing policies, assessing usage, and providing the information necessary to bill for services. These processes are considered important for effective network management and security.

This chapter includes the following sections:

- Authentication, page 32-1
- Authorization, page 32-2
- Accounting, page 32-2
- Interaction Between Authentication, Authorization, and Accounting, page 32-2
- AAA Servers, page 32-2
- AAA Server Groups, page 32-3
- Local Database Support, page 32-3
- Summary of AAA Service Support, page 32-3

Authentication

Authentication provides a way to identify a user, typically by having the user enter a valid username and valid password before access is granted. The AAA server compares a user's authentication credentials with other user credentials stored in a database. If the credentials match, the user is permitted access to the network. If the credentials do not match, authentication fails and network access is denied.

You can configure the ASA to authenticate the following items:

- All administrative connections to the ASA, including the following sessions:
  - Telnet
  - SSH. For more information, see Chapter 45, “Configuring Management Access.”
  - Serial console
  - ASDM using HTTPS
  - VPN management access
- The **enable** command. For more information, see Chapter 45, “Configuring Management Access.”

- Network access. For more information, see Chapter 38, "Configuring the Identity Firewall," Chapter 39, "Configuring the ASA to Integrate with Cisco TrustSec," and the firewall configuration guide.
- VPN access. For more information, see the VPN configuration guide.

Authorization

Authorization is the process of enforcing policies: determining what types of activities, resources, or services a user is permitted to access. After a user is authenticated, that user may be authorized for different types of access or activity.

You can configure the ASA to authorize the following items:

- Management commands. For more information, see Chapter 45, "Configuring Management Access."
- Network access. For more information, see the firewall configuration guide.
- VPN access. For more information, see the VPN configuration guide.

Accounting

Accounting measures the resources a user consumes during access, which may include the amount of system time or the amount of data that a user has sent or received during a session. Accounting is carried out through the logging of session statistics and usage information, which is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Interaction Between Authentication, Authorization, and Accounting

You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

AAA Servers

The AAA server is a network server that is used for access control. Authentication identifies the user. Authorization implements policies that determine which resources and services an authenticated user may access. Accounting keeps track of time and data resources that are used for billing and analysis.

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 or service.

Local Database Support

The ASA maintains a local database that you can populate with user profiles. You can use a local database instead of AAA servers to provide user authentication, authorization, and accounting. For more information, see Chapter 33, "Configuring the Local Database for AAA."

Summary of AAA Service Support

Table 32-1 provides cross-references to the configuration guide chapters that describe support for specific AAA service types.

Table 32-1 AAA Service Support

AAA Service	Configuration Guide Cross-Reference
Certificates	See Chapter 40, "Configuring Digital Certificates."
HTTP Form	See the VPN configuration guide.
Identity Firewall	See Chapter 38, "Configuring the Identity Firewall."
Kerberos	See the VPN configuration guide.
LDAP	See Chapter 36, "Configuring LDAP Servers for AAA."
Local Database	See Chapter 33, "Configuring the Local Database for AAA."
NT	See Chapter 37, "Configuring Windows NT Servers for AAA."
RADIUS	See Chapter 34, "Configuring RADIUS Servers for AAA."
RSA/SDI	See the VPN configuration guide.
TACACS+	See Chapter 35, "Configuring TACACS+ Servers for AAA."
TrustSec	See Chapter 39, "Configuring the ASA to Integrate with Cisco TrustSec."



Configuring the Local Database for AAA

This chapter describes how to configure local servers for AAAand includes the following sections:

- Information About the Local Database, page 33-1
- Fallback Support, page 33-2
- How Fallback Works with Multiple Servers in a Group, page 33-2
- Licensing Requirements for the Local Database, page 33-3
- Guidelines and Limitations, page 33-3
- Adding a User Account to the Local Database, page 33-3
- Testing Local Database Authentication and Authorization, page 33-7
- Monitoring the Local Database, page 33-7
- Feature History for the Local Database, page 33-8

Information About the Local Database

You can use the local database for the following functions:

- ASDM per-user access
- 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.

Note You cannot use the local database for network access authorization.

Fallback Support

The local database can act as a fallback method for several functions. This behavior is designed to help you prevent accidental lockout from the ASA.

When a user logs in, the servers in the group are accessed one at a time, starting with the first server that 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 have configured it as a fallback method (for management authentication and authorization only). If you do not have a fallback method, the ASA continues to try the AAA servers.

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.

Licensing Requirements for the Local Database

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

To prevent lockout from the ASA when using the local database for authentication or authorization, see the "Recovering from a Lockout" section on page 45-31.

Adding a User Account to the Local Database

To add a user to the local database, perform the following steps:

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 16-1), 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 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 6 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 45-17), 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 45-23 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 45-17), 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 45-17), 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 7 (Optional) To enable public key authentication for SSH connections to the ASA on a per-user basis, click one of the following options in the navigation pane:

- **Public Key Authentication**—Paste in a Base64-encoded public key. You can generate the key using any SSH key generation software (such as ssh keygen) that can generate SSH-RSA raw keys (with no certificates).When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, check the **Key is hashed** check box.

  To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.
- **Public Key Using PKF**—Check the **Specify a new PKF key** check box, and paste or import a public key file (PKF) formatted key, up to 4096 bits. Use this format for keys that are too large to paste in Base64 format. For example, you can generate a 4096-bit key using ssh keygen, then convert

it to PKF, and import on this pane. When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, copy it from the Public Key Athentication pane, and paste it in that pane on the new ASA with the **Key is hashed** check box checked.

To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.

Step 8 Click **VPN Policy** to configure VPN policy attributes for this user. See the Configuring VPN Policy Attributes for a Local User section in the VPN configuration guide.

Step 9 Click **Apply**.

The user is added to the local 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.

The following example generates a shared key for SSH on a Linux or Macintosh system, and imports it to the ASA:

Step 1 Generate the ssh-rsa public and private keys for 4096 bits on your computer:

```
jcrichton-mac:~ john$ ssh-keygen -b 4096
Generating public/private rsa key pair.
Enter file in which to save the key (/Users/john/.ssh/id_rsa):
/Users/john/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase): pa$$phrase
Enter same passphrase again: pa$$phrase
Your identification has been saved in /Users/john/.ssh/id_rsa.
Your public key has been saved in /Users/john/.ssh/id_rsa.pub.
The key fingerprint is:
c0:0a:a2:3c:99:fc:00:62:f1:ee:fa:f8:ef:70:c1:f9 john@jcrichton-mac
The key's randomart image is:
+--[ RSA 4096]----+
| .               |
|  o  .           |
|+...  o          |
|B.+.....         |
|.B ..+  S        |
|  =   o          |
|   + . E         |
|  o o            |
| ooooo           |
+-----------------+
```

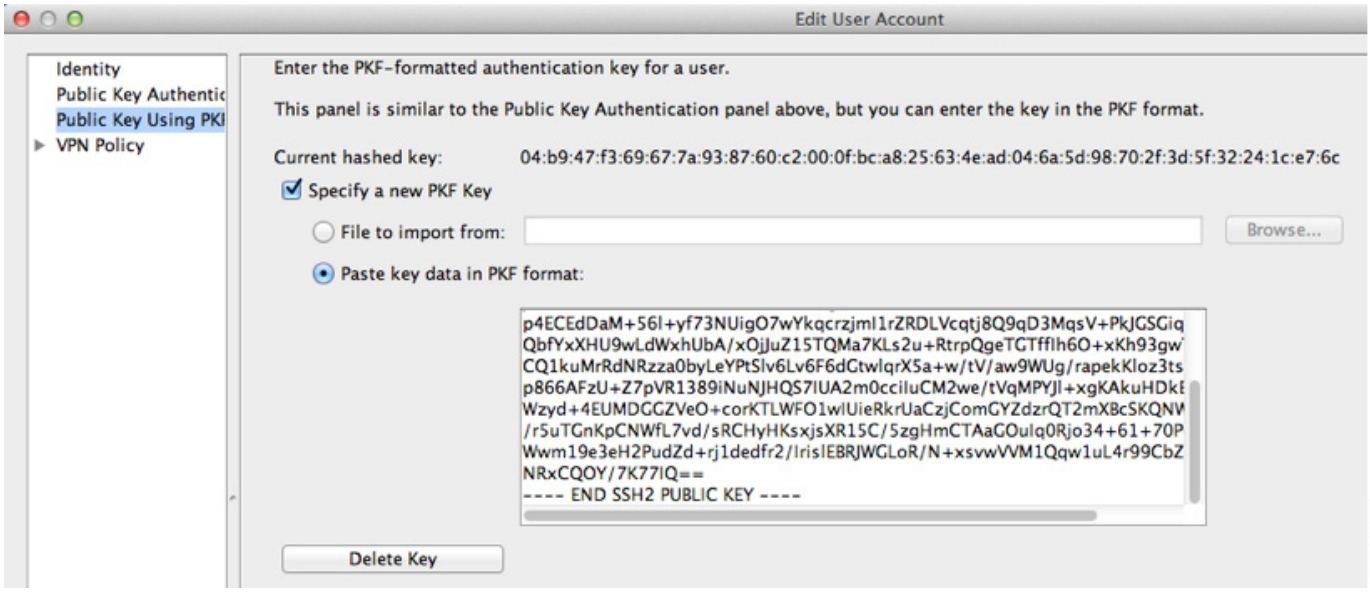
Step 2 Convert the key to PKF format:

```
jcrichton-mac:~ john$ cd .ssh
jcrichton-mac:.ssh john$ ssh-keygen -e -f id_rsa.pub
---- BEGIN SSH2 PUBLIC KEY ----
Comment: "4096-bit RSA, converted by ramona@rboersma-mac from OpenSSH"
AAAAB3NzaC1yc2EAAAADAQABAAACAQDNUvkgza37lB/Q/fljpLAv1BbyAd5PJCJXh/U4LO
hleR/qgIROjpnDaS7Az8/+sjHmq0qXC5TXkzWihvRZbhefyPhPHCi0hIt4oUF2ZbXESA/8
jUT4ehXIUE7FrChffBBtbD4d9FkV8A2gwZCDJBxEM26ocbZCSTx9QC//wt6E/zRcdoqiJG
```

```
p4ECEdDaM+56l+yf73NUigO7wYkqcrzjmI1rZRDLVcqtj8Q9qD3MqsV+PkJGSGiqZwnyIl
QbfYxXHU9wLdWxhUbA/xOjJuZ15TQMa7KLs2u+RtrpQgeTGTffIh6O+xKh93gwTgzaZTK4
CQ1kuMrRdNRzza0byLeYPtSlv6Lv6F6dGtwlqrX5a+w/tV/aw9WUg/rapekKloz3tsPTDe
p866AFzU+Z7pVR1389iNuNJHQS7IUA2m0cciIuCM2we/tVqMPYJl+xgKAkuHDkBlMS4i8b
Wzyd+4EUMDGGZVeO+corKTLWFO1wIUieRkrUaCzjComGYZdzrQT2mXBcSKQNWlSCBpCHsk
/r5uTGnKpCNWfL7vd/sRCHyHKsxjsXR15C/5zgHmCTAaGOuIq0Rjo34+61+70PCtYXebxM
Wwm19e3eH2PudZd+rj1dedfr2/IrislEBRJWGLoR/N+xsvwVVM1Qqw1uL4r99CbZF9NghY
NRxCQOY/7K77IQ==
---- END SSH2 PUBLIC KEY ----
jcrichton-mac:.ssh john$
```

Step 3 Copy the key to your clipboard.

Step 4 In ASDM, choose **Configuration > Device Management > Users/AAA > User Accounts**, select the username and then click **Edit**. Click **Public Key Using PKF** and paste the key into the window:



Step 5 Verify the user (test) can SSH to the ASA:

```
jcrichton-mac:.ssh john$ ssh test@10.86.118.5
The authenticity of host '10.86.118.5 (10.86.118.5)' can't be established.
RSA key fingerprint is 39:ca:ed:a8:75:5b:cc:8e:e2:1d:96:2b:93:b5:69:94.
Are you sure you want to continue connecting (yes/no)? yes
```

The following dialog box appears for you to enter your passphrase:



Meanwhile, in the terminal session:

```
Warning: Permanently added '10.86.118.5' (RSA) to the list of known hosts.
Identity added: /Users/john/.ssh/id_rsa (/Users/john/.ssh/id_rsa)
Type help or '?' for a list of available commands.
asa>
```

Testing Local Database Authentication and Authorization

To determine whether the ASA can contact a local database 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 the Local Database

To monitor the local database, see the following panes:

Path	Purpose
Monitoring > Properties > AAA Servers	Shows the configured database statistics.
Monitoring > Properties > AAA Servers	Shows the AAA server running configuration.

Feature History for the Local Database

Table 33-1 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 33-1 Feature History for the Local Database

Feature Name	Platform Releases	Feature Information
Local database configuration for AAA	7.0(1)	Describes how to configure the local database for AAA use. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > User Accounts.
Support for SSH public key authentication	9.1(2)	You can now enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF *Also available in 8.4(4.1); PKF key format support is only in 9.1(2).*



Configuring RADIUS Servers for AAA

This chapter describes how to configure RADIUS servers for AAA and includes the following sections:

- Information About RADIUS Servers, page 34-1
- Licensing Requirements for RADIUS Servers, page 34-13
- Guidelines and Limitations, page 34-14
- Configuring RADIUS Servers, page 34-14
- Testing RADIUS Server Authentication and Authorization, page 34-18
- Monitoring RADIUS Servers, page 34-19
- Additional References, page 34-20
- Feature History for RADIUS Servers, page 34-20

Information About RADIUS Servers

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

This section includes the following topics:

- Supported Authentication Methods, page 34-2
- User Authorization of VPN Connections, page 34-2
- Supported Sets of RADIUS Attributes, page 34-2
- Supported RADIUS Authorization Attributes, page 34-3
- Supported IETF RADIUS Authorization Attributes, page 34-12
- RADIUS Accounting Disconnect Reason Codes, page 34-13

Supported Authentication Methods

The ASA supports the following authentication methods with RADIUS servers:

- 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—For 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.

User Authorization of VPN Connections

The ASA can use RADIUS servers for user authorization of VPN remote access and firewall cut-through-proxy sessions using dynamic ACLs or ACL names per user. To implement dynamic ACLs, you must configure the RADIUS server to support them. When the user authenticates, the RADIUS server sends a downloadable ACL or ACL name to the ASA. Access to a given service is either permitted or denied by the ACL. The ASA deletes the ACL when the authentication session expires.

In addition to ACLs, the ASA supports many other attributes for authorization and setting of permissions for VPN remote access and firewall cut-through proxy sessions.

Supported Sets of RADIUS Attributes

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.

Supported 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 34-1 lists the 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.

All attributes listed in Table 34-1 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 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 Version 9.0(1).

Table 34-1 Supported RADIUS Authorization Attributes

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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
Group-Policy	Y	25	String	Single	Sets the group policy for the remote access VPN session. For Versions 8.2.x and later, use this attribute instead of IETF-Radius-Class. You can use one of the 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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
IPsec-Over-UDP	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP-Port	Y	35	Integer	Single	4001- 49151. The default is 10000.
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 or 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
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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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)
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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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.
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

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
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	0-30. 0 = Disabled.
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://www.cisco.com/en/US/docs/security/asa/asa80/asdm60/ssl_vpn_deployment_guide/deploy.html
WebVPN-Macro-Substitution	Y	224	String	Single	Unbounded. For examples, see the *SSL VPN Deployment Guide* at the following URL: http://www.cisco.com/en/US/docs/security/asa/asa80/asdm60/ssl_vpn_deployment_guide/deploy.html
WebVPN-Port-Forwarding-Enable	Y	97	Integer	Single	0 = Disabled 1 = Enabled

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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	0-30. 0 = Disabled.
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
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.

Table 34-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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

Supported IETF RADIUS Authorization Attributes

Table 34-2 lists the supported IETF RADIUS attributes.

Table 34-2 Supported IETF RADIUS Attributes

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
IETF-Radius-Class	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 34-1: • *group policy name* • OU=*group policy name* • OU=*group policy name*
IETF-Radius-Filter-Id	Y	11	String	Single	ACL name that is defined on the ASA, which applies only to full tunnel IPsec and SSL VPN clients.
IETF-Radius-Framed-IP-Address	Y	n/a	String	Single	An IP address
IETF-Radius-Framed-IP-Netmask	Y	n/a	String	Single	An IP address mask
IETF-Radius-Idle-Timeout	Y	28	Integer	Single	Seconds

Table 34-2 Supported IETF RADIUS Attributes (continued)

IETF-Radius-Service-Type	Y	6	Integer	Single	Seconds. Possible Service Type values: • .Administrative—User is allowed access to the configure prompt. • .NAS-Prompt—User is allowed access to the exec prompt. • .remote-access—User is allowed network access
IETF-Radius-Session-Timeout	Y	27	Integer	Single	Seconds

RADIUS Accounting Disconnect Reason Codes

These codes are returned if the ASA encounters a disconnect when sending packets:

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
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

Licensing Requirements for RADIUS 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.

Additional 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.
- If you need to configure fallback support using the local database, see the "Fallback Support" section on page 33-2 and the "How Fallback Works with Multiple Servers in a Group" section on page 33-2.
- To prevent lockout from the ASA when using RADIUS authentication, see the "Recovering from a Lockout" section on page 45-31.

Configuring RADIUS Servers

This section includes the following topics:

- Task Flow for Configuring RADIUS Servers, page 34-14
- Configuring RADIUS Server Groups, page 34-15
- Adding a RADIUS Server to a Group, page 34-16
- Adding an Authentication Prompt, page 34-18

Task Flow for Configuring RADIUS Servers

Step 1 Load the ASA attributes into the RADIUS server. The method that you use to load the attributes depends on which type of RADIUS server that 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).

Step 2 Add a RADIUS server group. See the "Configuring RADIUS Server Groups" section on page 34-15.

Step 3 For a server group, add a server to the group. See the "Adding a RADIUS Server to a Group" section on page 34-16.

Step 4 (Optional) Specify text to display to the user during the AAA authentication challenge process. See the "Adding an Authentication Prompt" section on page 34-18.

Configuring RADIUS Server Groups

If you want to use an external RADIUS server for authentication, authorization, or accounting, you must first create at least one RADIUS server group per AAA protocol and add one or more servers to each group. You identify AAA server groups by name.

To add a RADIUS server group, perform the following steps:

Detailed 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 RADIUS server type.

Step 5 In the Accounting Mode field, click **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.

Step 6 In the Reactivation Mode field, click **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 options 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 to the running configuration.

Adding a RADIUS Server to a Group

To add a RADIUS 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 length 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 In the ACL Netmask Convert field, specify how you want the ASA to handle netmasks received in downloadable ACLs. Choose from the following options:

- 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 ACLs received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed.
- Wildcard—The ASA assumes downloadable ACLs received from the RADIUS server contain only wildcard netmask expressions, and it converts them all to standard netmask expressions when the ACLs are downloaded.

Step 7 In the Common Password field, specify 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 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.

Although the password is required by the RADIUS protocol and the RADIUS server, users do not need to know it.

Step 8 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.

Step 9 In the Retry Interval field, specify the length of time, from 1 to 10 seconds, that the ASA waits between attempts to contact the server.

Step 10 In the Accounting Mode field, click **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.

Step 11 In the Server Accounting Port field, specify the server port to be used for accounting of users. The default port is 1646.

Step 12 In the Server Authentication Port field, specify the server port to be used for authentication of users. The default port is 1645.

Step 13 In the Server Secret Key field, specify the shared secret key used to authenticate the RADIUS server to the ASA. The server secret that you configure 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.

Step 14 Click **OK**.

The Add AAA Server Group dialog box closes, and the AAA server is added to the AAA server group.

Step 15 In the AAA Server Groups pane, click **Apply** to save the changes to the running configuration.

Adding an Authentication Prompt

You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from 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 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
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 either accepted or rejected by the RADIUS server.

If the RADIUS 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** to save the changes to the running configuration.

Testing RADIUS Server Authentication and Authorization

To determine whether the ASA can contact a RADIUS 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 RADIUS Servers

To monitor RADIUS servers, see the following panes:

Path	Purpose
Monitoring > Properties > AAA Servers	Shows the configured RADIUS server statistics.
Monitoring > Properties > AAA Servers	Shows the RADIUS server running configuration.

Additional References

For additional information related to implementing AAA through RADIUS servers, see the "RFCs" section on page 34-20.

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 RADIUS Servers

Table 34-3 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 34-3 Feature History for RADIUS Servers

Feature Name	Platform Releases	Feature Information
RADIUS Servers for AAA	7.0(1)	Describes how to configure RADIUS servers for AAA. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups 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 TACACS+ Servers for AAA

This chapter describes how to configure TACACS+ servers used in AAA and includes the following sections:

- Information About TACACS+ Servers, page 35-1
- Licensing Requirements for TACACS+ Servers, page 35-2
- Guidelines and Limitations, page 35-3
- Configuring TACACS+ Servers, page 35-3
- Testing TACACS+ Server Authentication and Authorization, page 35-6
- Monitoring TACACS+ Servers, page 35-7
- Feature History for TACACS+ Servers, page 35-7

Information About TACACS+ Servers

The ASA supports TACACS+ server authentication with the following protocols: ASCII, PAP, CHAP, and MS-CHAPv1.

Using TACACS+ Attributes

The ASA provides support for TACACS+ attributes. TACACS+ attributes separate 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 35-1 lists supported TACACS+ authorization response attributes for cut-through-proxy connections. Table 35-2 lists supported TACACS+ accounting attributes.

Table 35-1 Supported TACACS+ Authorization Response Attributes

Attribute	Description
acl	Identifies a locally configured ACL 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 35-2 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.

Licensing Requirements for TACACS+ 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.

Additional 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.
- If you need to configure fallback support using the local database, see the "Fallback Support" section on page 33-2 and the "How Fallback Works with Multiple Servers in a Group" section on page 33-2.
- To prevent lockout from the ASA when using TACACS+ authentication or authorization, see the "Recovering from a Lockout" section on page 45-31.

Configuring TACACS+ Servers

This section includes the following topics:

- Task Flow for Configuring TACACS+ Servers, page 35-3
- Configuring TACACS+ Server Groups, page 35-4
- Adding a TACACS+ Server to a Group, page 35-4
- Adding an Authentication Prompt, page 35-5

Task Flow for Configuring TACACS+ Servers

Step 1 Add a TACACS+ server group. See the "Configuring TACACS+ Server Groups" section on page 35-4.

Step 2 For a server group, add a server to the group. See the "Adding a TACACS+ Server to a Group" section on page 35-4.

Step 3 (Optional) Specify text to display to the user during the AAA authentication challenge process. See the "Adding an Authentication Prompt" section on page 35-5.

Configuring TACACS+ Server Groups

If you want to use a TACACS+ server for authentication, authorization, or accounting, you must first create at least one TACACS+ server group and add one or more servers to each group. You identify TACACS+ server groups by name.

To add a TACACS+ server group, perform the following steps:

Detailed 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 TACACS+ server type:

Step 5 In the Accounting Mode field, click **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.

Step 6 In the Reactivation Mode field, click **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 Click **OK**.

The Add AAA Server Group dialog box closes, and the new server group is added to the AAA Server Groups table.

Step 10 In the AAA Server Groups dialog box, click **Apply** to save the changes to the running configuration.

Adding a TACACS+ Server to a Group

To add a TACACS+ 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 Specify the server port. The server port is either port number 139, or the TCP port number used by the ASA to communicate with the TACACS+ server.

Step 7 Specify the 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 TACACS+ server administrator. The maximum field length is 64 characters.

Step 8 Click **OK**.

The Add AAA Server Group dialog box closes, and the AAA server is added to the AAA server group.

Step 9 In the AAA Server Groups pane, click **Apply** to save the changes to the running configuration.

Adding an Authentication Prompt

You can specify text to display to the user during the AAA authentication challenge process. You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from TACACS+ 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 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 Choose **Configuration > Device Management > Users/AAA > Authentication Prompt.**

Step 2 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 3 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 4 Click **Apply** to save the changes to the running configuration.

Testing TACACS+ Server Authentication and Authorization

To determine whether the ASA can contact a TACACS+ 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 TACACS+ Servers

To monitor TACACS+ servers, see the following panes:

Path	Purpose
Monitoring > Properties > AAA Servers	Shows the configured TACACS+ server statistics.
Monitoring > Properties > AAA Servers	Shows the TACACS+ server running configuration.

Feature History for TACACS+ Servers

Table 35-3 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 35-3 Feature History for TACACS+ Servers

Feature Name	Platform Releases	Feature Information
TACACS+ Servers	7.0(1)	Describes how to configure TACACS+ servers for AAA. , We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > Authentication Prompt.



Configuring LDAP Servers for AAA

This chapter describes how to configure LDAP servers used in AAA and includes the following sections:

- Information About LDAP and the ASA, page 36-1
- Licensing Requirements for LDAP Servers, page 36-4
- Guidelines and Limitations, page 36-4
- Configuring LDAP Servers, page 36-5
- Testing LDAP Server Authentication and Authorization, page 36-9
- Monitoring LDAP Servers, page 36-10
- Feature History for LDAP Servers, page 36-10

Information About LDAP and the ASA

The ASA is compatible with the most LDAPv3 directory servers, including:

- Sun Microsystems JAVA System Directory Server, now part of Oracle Directory Server Enterprise Edition, and formerly named the Sun ONE Directory Server
- Microsoft Active Directory
- Novell
- OpenLDAP

By default, the ASA autodetects whether it is connected to Microsoft Active Directory, Sun LDAP, Novell, OpenLDAP, or a generic LDAPv3 directory server. However, if autodetection fails to determine the LDAP server type, you can manually configure it.

LDAP Server Guidelines

When configuring the LDAP server, 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 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.

How Authentication Works 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.

The ASA and LDAP server supports 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 one 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 two.

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. In this case, using LDAP accomplishes authentication and authorization in a single step.

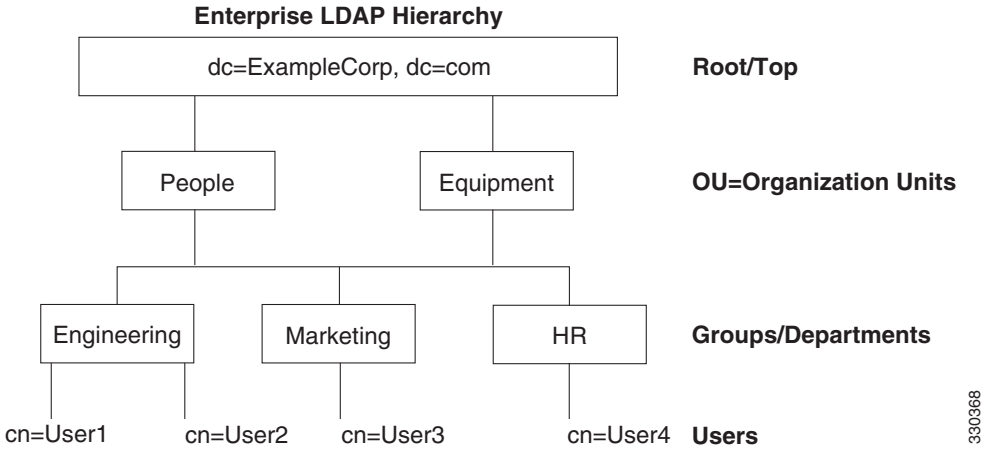
Note For more information about the LDAP protocol, see RFCs 1777, 2251, and 2849.

About the LDAP Hierarchy

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 36-1 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 36-1 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 you are looking for. Together, these fields 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 36-1 shows a sample LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. Table 36-1 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 36-1 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

About Binding to an LDAP Server

The ASA uses the login DN and login password to establish trust (bind) with an LDAP server. When performing a Microsoft Active Directory read-only operation (such as authentication, authorization, or group search), the ASA can bind using 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 the following authentication methods:

- 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.

Note As an LDAP client, the ASA does not support the transmission of anonymous binds or requests.

Licensing Requirements for LDAP 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 LDAP Servers

This section includes the following topics:

- Task Flow for Configuring LDAP Servers, page 36-5
- Configuring LDAP Attribute Maps, page 36-5
- Configuring LDAP Server Groups, page 36-7
- Adding an LDAP Server to a Group, page 36-8

Task Flow for Configuring LDAP Servers

Step 1 Add an LDAP server group. See the "Configuring LDAP Server Groups" section on page 36-7.

Step 2 Add a server to the group, then configure server parameters. See the "Adding an LDAP Server to a Group" section on page 36-8.

Step 3 Configure LDAP attribute maps. See the "Configuring LDAP Attribute Maps" section on page 36-5. You must add an attribute map before adding an LDAP server to an LDAP server group.

Configuring LDAP Attribute Maps

The ASA can use an LDAP directory for authenticating users for:

- VPN remote access users
- firewall network access/cut-through-proxy sessions
- setting policy permissions (also called authorization attributes), such as ACLs, bookmark lists, DNS or WINS settings, and session timers.
- setting the key attributes in a local group policy

The ASA uses LDAP attribute maps to translate native LDAP user attributes to Cisco ASA attributes. You can bind these attribute maps to LDAP servers or remove them. You can also show or clear attribute maps.

Guidelines

The LDAP attribute map does not support multi-valued attributes. For example, if a user is a member of several AD groups, and the LDAP attribute map matches more than one group, the value chosen is based on the alphabetization of the matched entries.

To use the attribute mapping features correctly, you need to understand LDAP attribute names and values, as well as the user-defined attribute names and values.

The names of frequently mapped 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 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—Applies an access control list or ACL 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 map one LDAP attribute from a specific LDAP server.

To map LDAP features, perform the following steps:

Detailed Steps

Step 1 Choose **Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map** (for local users), or **Configuration > Device Management > Users/AAA > LDAP Attribute Map (for all other users)**, then click **Add**.

The Add LDAP Attribute Map dialog box appears with the Map Name tab active.

Step 2 In the Name field, create a name for this attribute map.

Step 3 In the LDAP Attribute Name field, add the name of on of the LDAP attributes to be mapped.

Step 4 From the Cisco Attribute Name drop-down list, choose a Cisco attribute.

Step 5 Click **Add**.

Step 6 The attributes are mapped. To map more attributes, repeat Steps 1 through 5.

Step 7 If you want to map the value of any of the LDAP attributes to a new value in the mapped Cisco attribute, click the **Map Value** tab.

Step 8 Click **Add**.

The Add Mapping of Attribute Name dialog box appears.

Step 9 Choose an LDAP attribute from the drop-down list.

Step 10 In the LDAP Attribute Value field, enter the value for this LDAP attribute that you expect to be returned from the LDAP server.

Step 11 In the Cisco Attribute Value field, enter the value you want to use in the Cisco attribute when this LDAP attribute contains the previous LDAP Attribute Value.

Step 12 Click **Add**.

The values are mapped.

Step 13 To map more values, 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** to save the value mappings to the running configuration.

Configuring LDAP Server Groups

To use an external LDAP server for authentication, authorization, and/or accounting, you must first create at least one LDAP server group, and add one or more servers to each group. You identify LDAP server groups by name. Each server group is specific to one type of server.

Guidelines

- You can have up to 100 LDAP server groups in single mode or 4 LDAP server groups per context in multiple mode.
- Each group can have up to 16 LDAP servers in single mode or 4 LDAP servers in multiple mode.
- When a user logs in, the LDAP servers are accessed one at a time, starting with the first server that 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 LDAP servers.

Detailed Steps

The following steps show how to create and configure an LDAP server group, and add an LDAP server to that group.

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups,** or **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups** for VPN users.

Step 2 In the AAA Server Groups area, click **Add**.

The Add AAA Server Group dialog box appears.

Step 3 In the AAA Server Group field, name this AAA server group.

Step 4 From the Protocol drop-down list, choose the LDAP server type.

Step 5 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.

a. 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 6 In the Max Failed Attempts field, add the number of failed attempts to connect to the server to allow.

This option sets the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive.

Step 7 Click **OK**.

The Add AAA Server Group dialog box closes, and the new server group is added to the AAA Server Groups table.

Step 8 In the AAA Server Groups dialog box, click **Apply** to save the changes.

The changes are saved to the running configuration.

Adding an LDAP Server to a Group

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**, or **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups** for VPN users, and in the AAA Server Groups area, select the server group to which you want to add a server.

Step 2 Next to the list of Servers in Selected Grous, click **Add**.

The Add AAA Server dialog box appears for the selected server group.

Step 3 From the Interface Name drop-down list, choose the name of the interface that connects to the LDAP server.

Step 4 In the Server Name or IP Address field, add either the server name or IP address of the LDAP server.

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 In the LDAP Parameters for authentication/authorization area, configure the following fields:

- **Enable LDAP over SSL** (also called secure LDAP or LDAP-S)—Check this if you want to use SSL to secure communications between the ASA and the LDAP server.

Note If you do not configure the SASL protocol, we strongly recommend that you secure LDAP communications with SSL.

- **Server Port**—Enter 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.
- **Server Type**—Specify the LDAP server type from the drop-down list. The available options include the following:
  - Detect Automatically/Use Generic Type
  - Microsoft
  - Novell
  - OpenLDAP
  - Sun, now part of Oracle Directory Server Enterprise Edition
- **Base DN**—Enter 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**—Specify the extent of the search that the server should perform in the LDAP hierarchy when it receives an authorization request from the drop-down list. The following options are available:
  - 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)**—Enter the Relative Distinguished Name attribute(s) that uniquely identify an entry on the LDAP server. Common naming attributes are Common Name (CN), sAMAccountName, userPrincipalName, and User ID (uid).
- **Login DN and Login Password**—The ASA uses the login DN and login password to establish trust (bind) with an LDAP server. Specify the login password, which is the password for the login DN user account. The characters that you type are replaced with asterisks.
- **LDAP Attribute Map**—Select one of the attribute maps that you created for this LDAP server to use. These attribute maps map LDAP attribute names to Cisco attribute names and values.
- **SASL MD5 authentication** This enables the MD5 mechanism of the SASL to authenticate communications between the ASA and the LDAP server.
- **SASL Kerberos authentication**—Enables the Kerberos mechanism of the SASL to secure authentication communications between the ASA and the LDAP server. You must have defined a Kerberos server in order to enable this option.
- **LDAP Parameters for Group Search**—The fields in this area configure how the ASA requests AD groups.
  - **Group Base 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**—Specifiy the maximum time to wait for a response from an AD server that was queried for available groups.

Step 7 Click **OK**.

The Add AAA Server 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 to the running configuration.

Testing LDAP Server Authentication and Authorization

To determine whether the ASA can contact an LDAP server and authenticate or authorize a user, perform the following steps:

Step 1 In the Configuration > Device Management > Users/AAA > AAA Server Groups pane, select the server group in which the server resides.

Step 2 From the Servers in the Selected Group area, select the server that you want to test.

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 Enter a username.

Step 6 If you are testing authentication, enter the password for the username.

Step 7 Click **OK**.

The ASA sends either an authentication or authorization test message to the server. If the test fails, ASDM displays an error message.

Monitoring LDAP Servers

To monitor LDAP servers, perform the following steps:

Step 1 In ASDM, choose **Monitoring > Properties > AAA Servers**.

Step 2 To update an LDAP server status, select it then click **Update Server Statistics**.

The Update AAA Server Status dialog box appears with the LDAP server selected in the drop-down list.

Step 3 Click **OK**.

Step 4 To update the currently displayed statistics, click **Clear Server Statistics**.

Feature History for LDAP Servers

Table 36-2 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 36-2 Feature History for AAA Servers

Feature Name	Platform Releases	Feature Information
LDAP Servers for AAA	7.0(1)	LDAP Servers describe support for AAA and how to configure LDAP servers. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map



Configuring Windows NT Servers for AAA

This chapter describes how to configure Windows NT servers used in AAA and includes the following sections:

- Information About Windows NT Servers, page 37-1
- Licensing Requirements for Windows NT Servers, page 37-1
- Guidelines and Limitations, page 37-2
- Configuring Windows NT Servers, page 37-2
- Testing Windows NT Server Authentication and Authorization, page 37-4
- Monitoring Windows NT Servers, page 37-5
- Feature History for Windows NT Servers, page 37-5

Information About Windows NT Servers

The ASA supports Microsoft Windows server operating systems that support NTLM Version 1, collectively referred to as NT servers.

Note Windows NT servers have a maximum length of 14 characters for user passwords. Longer passwords are truncated, which is a limitation of NTLM Version 1.

Licensing Requirements for Windows NT 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.

Additional 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.
- If you need to configure fallback support using the local database, see the "Fallback Support" section on page 33-2 and the "How Fallback Works with Multiple Servers in a Group" section on page 33-2.

Configuring Windows NT Servers

This section includes the following topics:

- Configuring Windows NT Server Groups, page 37-3
- Adding a Windows NT Server to a Group, page 37-3

Task Flow for Configuring Windows NT Servers

Step 1 Add a AAA server group. See the "Configuring Windows NT Server Groups" section on page 37-3.

Step 2 For a server group, add a server to the group. See the "Adding a Windows NT Server to a Group" section on page 37-3.

Configuring Windows NT Server Groups

If you want to use a Windows NT server for authentication, authorization, or accounting, you must first create at least one Windows NT server group and add one or more servers to each group. You identify Windows NT server groups by name.

To add a Windows NT server group, perform the following steps:

Detailed 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 NT Domain server type.

Step 5 In the Reactivation Mode field, click **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 6 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 7 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 8 Click **OK**.

The Add AAA Server Group dialog box closes, and the new server group is added to the AAA Server Groups table.

Step 9 In the AAA Server Groups dialog box, click **Apply** to save the changes to the running configuration.

Adding a Windows NT Server to a Group

To add a Windows NT 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 Specify the server port. The server port is either port number 139, or the TCP port number used by the ASA to communicate with the Windows NT server.

Step 7 Specify the name of the domain controller. The domain controller is the hostname (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 hostname for the server whose IP address you added in the Authentication Server Address field. If the name is incorrect, authentication fails.

Step 8 Click **OK**.

The Add AAA Server Group dialog box closes, and the AAA server is added to the AAA server group.

Step 9 In the AAA Server Groups pane, click **Apply** to save the changes to the running configuration.

Testing Windows NT Server Authentication and Authorization

To determine whether the ASA can contact a Windows NT server and authenticate or authorize a user, perform the following steps:

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**, and in the AAA Server Groups area, click the server group 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 Windows NT Servers

To monitor Windows NT servers, see the following panes:

Path	Purpose
Monitoring > Properties > AAA Servers	Shows the configured Windows NT server statistics.
Monitoring > Properties > AAA Servers	Shows the Windows NT server running configuration.

Feature History for Windows NT Servers

Table 37-1 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 37-1 Feature History for Windows NT Servers

Feature Name	Platform Releases	Feature Information
Windows NT Servers for AAA	7.0(1)	Describes support for Windows NT Servers and how to configure them for AAA. We introduced the following screen: Configuration > Device Management > Users/AAA > AAA Server Groups.



Configuring the Identity Firewall

This chapter describes how to configure the ASA for the Identity Firewall and includes the following sections:

- Information About the Identity Firewall, page 38-1
- Licensing for the Identity Firewall, page 38-7
- Guidelines and Limitations, page 38-8
- Prerequisites, page 38-9
- Configuring the Identity Firewall, page 38-10
- Monitoring the Identity Firewall, page 38-17
- Feature History for the Identity Firewall, page 38-19

Information About the Identity Firewall

This section includes the following topics:

- Overview of the Identity Firewall, page 38-1
- Architecture for Identity Firewall Deployments, page 38-2
- Features of the Identity Firewall, page 38-3
- Deployment Scenarios, page 38-4

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 username/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 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 usernames 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
- Simplifying 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

  Although Active Directory is part of the Identity Firewall on the ASA, Active Directory administrators manage it. The reliability and accuracy of the data depends on data in Active Directory.

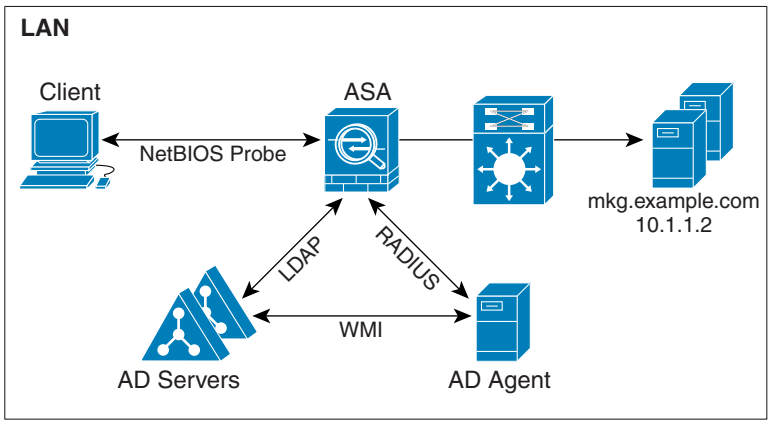
  Supported versions include Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 servers.
- Active Directory (AD) Agent

  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 38-1 show the components of the Identity Firewall. The succeeding table describes the roles of these components and how they communicate with one another.

Figure 38-1 Identity Firewall Components



1	**On the ASA**: Administrators configure local user groups and Identity Firewall policies.	**4**	**Client <-> ASA**: The client logs into the network through Microsoft Active Directory. The AD Server authenticates users and generates user login security logs. Alternatively, the client can log into the network through a cut-through proxy or 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 Modular Policy Framework 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 that asks for the user's IP address. The ASA forwards the new mapped entries that have been learned from web authentication and VPN sessions to the AD Agent.	**6**	**AD Agent <-> AD Server**: The AD Agent maintains a cache of user ID and IP address mapped entries. and notifies the ASA of changes. The AD Agent sends logs to a syslog server.

Features of the Identity Firewall

The Identity Firewall includes the following key features.

Flexibility

- The ASA can retrieve user identity and IP address mapping 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 the existing 5-tuple solution.
- Supports use 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 ASAs that are connected to the AD Agent.

Scalability

- Each AD Agent supports 100 ASAs. Multiple ASAs 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 mapped entries in active policies for the ASA 5500 Series models. This limit controls the maximum number of users who have policies applied. The total number of users are the aggregate of all users configured in all different contexts.
- Supports up to 1024 user identity-IP address mapped entries in active policies for the ASA 5505.
- Supports up to 256 user groups in active ASA policies.
- A single access rule can contain one or more user groups or users.
- Supports multiple domains.

Availability

- The ASA retrieves group information from the Active Directory and falls back to web authentication for IP addresses when 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 for use among ASAs.

Deployment Scenarios

You can deploy the components of the Identity Firewall in the following ways, depending on your environmental requirements.

Figure 38-2 shows how 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 this deployment scenario, the Active Directory server and AD Agent are co-located on the same Windows server.

Figure 38-2 Deployment Scenario without Redundancy

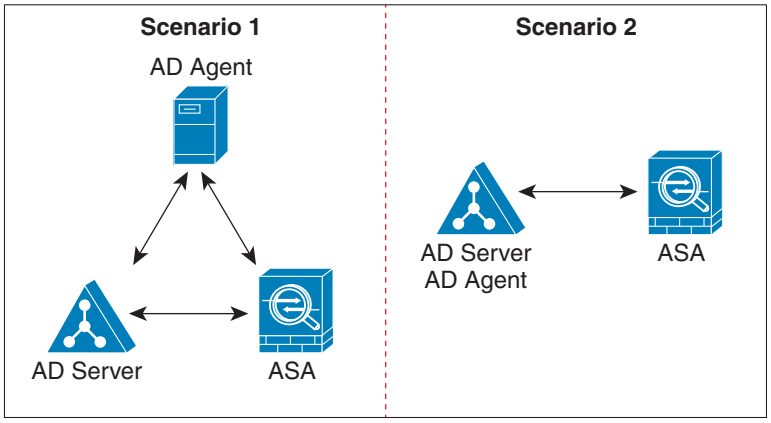


Figure 38-3 shows how you can deploy the Identity Firewall components to support redundancy. Scenario 1 shows a deployment with multiple Active Directory servers and a single AD Agent installed on a separate Windows server. Scenario 2 shows a deployment with multiple Active Directory servers and multiple AD Agents installed on separate Windows servers.

Figure 38-3 Deployment Scenario with Redundant Components

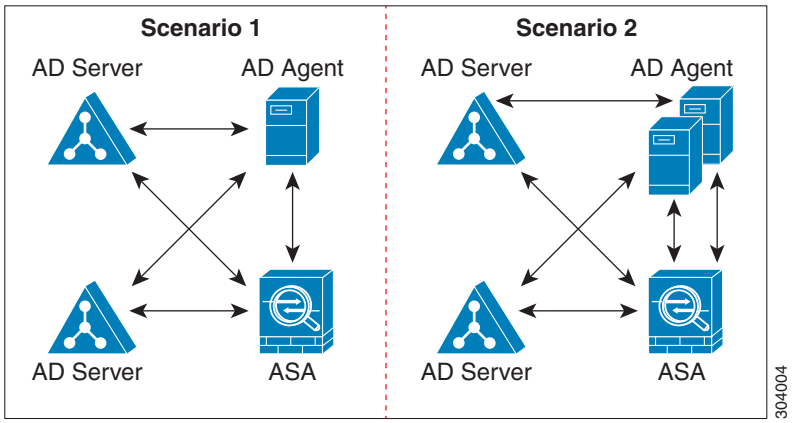


Figure 38-4 shows how all Identity Firewall components—Active Directory server, the AD Agent, and the clients—are installed and communicate on the LAN.

Figure 38-4 LAN -based Deployment

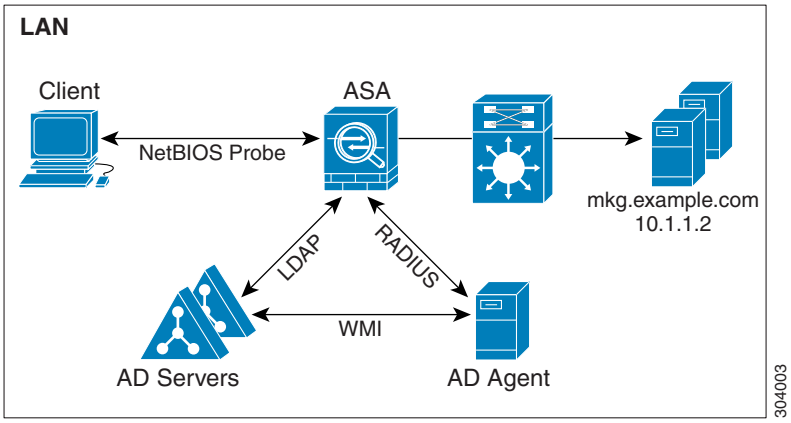


Figure 38-5 shows a WAN-based deployment to support a remote site. The Active Directory server and the AD Agent are installed on the main site LAN. The clients are located at a remote site and connect to the Identity Firewall components over a WAN.

Figure 38-5 WAN-based Deployment

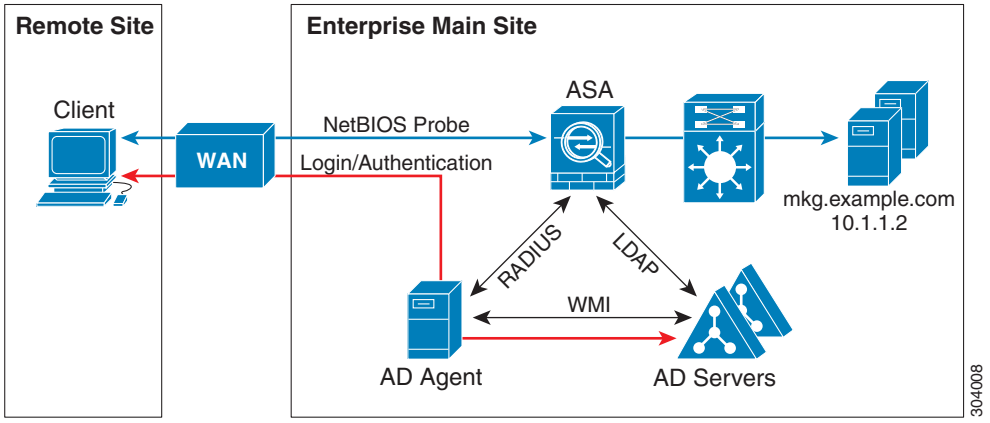


Figure 38-6 also shows a WAN-based deployment to support a remote site. The Active Directory server is installed on the main site LAN. However, the AD Agent is installed and accessed by the clients at the remote site. The remote clients connect to the Active Directory servers at the main site over a WAN.

Figure 38-6 WAN-based Deployment with Remote AD Agent

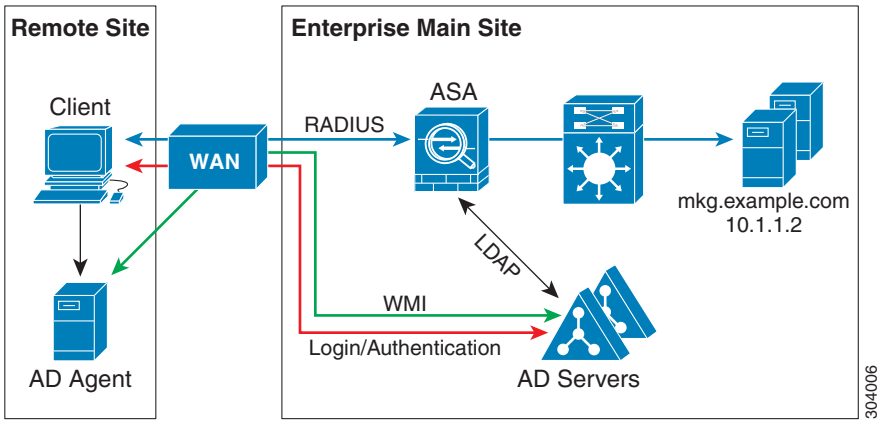
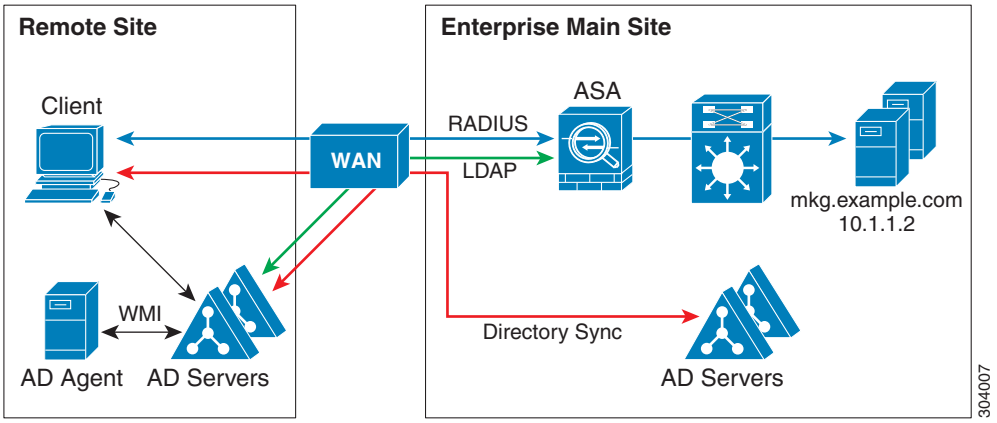


Figure 38-7 shows an expanded remote site installation. An AD Agent and Active Directory servers are installed at the remote site. The clients access these components locally when logging into network resources located at the main site. The remote Active Directory server must synchronize its data with the central Active Directory servers located at the main site.

Figure 38-7 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 mapping and AD Agent status replication from active to standby when Stateful Failover is enabled. However, only user identity-IP address mapping, 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 to be 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 ASAs. 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 into 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 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 maps
  - Crypto maps
  - WCCP
  - NAT
  - Group policy (except for VPN filters)

  - DAP

- You can use the **user-identity update active-user-database** command to actively initiate a user-IP address download from the AD agent.

  By design, if a previous download session has finished, the ASA does not allow you to issue this command again.

  As a result, if the user-IP database is very large, the previous download session is not finished yet, and you issue another **user-identity update active-user-database** command, the following error message appears:

  ```
  "ERROR: one update active-user-database is already in progress."
  ```

  You need to wait until the previous session is completely finished, then you can issue another **user-identity update active-user-database** command.

  Another example of this behavior occurs because of packet loss from the AD Agent to the ASA.

  When you issue a **user-identity update active-user-database** command, the ASA requests the total number of user-IP mapped entries to be downloaded. Then the AD Agent initiates a UDP connection to the ASA and sends the change of authorization request packet.

  If for some reason the packet is lost, there is no way for the ASA to discern this. As a result, the ASA holds the session for 4-5 minutes, during which time this error message continues to appear if you have issued the **user-identity update active-user-database** command.

- When you use the Cisco Context Directory Agent (CDA) in conjunction with the ASA or Cisco Ironport Web Security Appliance (WSA), make sure that you open the following ports:
  - Authentication port for UDP—1645
  - Accounting port for UDP—1646
  - Listening port for UDP—3799

    The listening port is used to send change of authorization requests from the CDA to the ASA or to the WSA.
- For domain names, the following characters are not valid: \/:*?"<>|. For naming conventions, see http://support.microsoft.com/kb/909264.
- For usernames, the following characters are not valid: \/[]:;=,+*?"<>|@.
- For user group names, the following characters are not valid: \/[]:;=,+*?"<>|.

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 and 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 Microsoft Active Directory for how to enable SSL for Active Directory.

Note Before running the AD Agent Installer, you must install the patches listed in the *README First for the Cisco Active Directory Agent* on each 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.

Configuring the Identity Firewall

This section contains the following topic:

- Task Flow for Configuring the Identity Firewall, page 38-10
- Configuring the Active Directory Domain, page 38-11
- Configuring Active Directory Server Groups, page 38-12
- Configuring Active Directory Agents, page 38-12
- Configuring Active Directory Agent Groups, page 38-13
- Configuring Identity Options, page 38-13
- Configuring Identity-Based Security Policy, page 38-16

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 the "Configuring the Active Directory Domain" section on page 38-11 and the "Configuring Active Directory Server Groups" section on page 38-12.

See also the "Deployment Scenarios" section on page 38-4 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 the "Configuring Active Directory Server Groups" section on page 38-12 and the "Configuring Active Directory Agent Groups" section on page 38-13.

See also "Deployment Scenarios" section on page 38-4 for the ways in which you can deploy the AD Agents to meet your environment requirements.

Step 3 Configure Identity Options.

See the "Configuring Identity Options" section on page 38-13.

Step 4 Configure Identity-based Security Policy. After the AD domain and AD Agent are configured, you can create identity-based object groups and ACLs for use in many features.

See the "Configuring Identity-Based Security Policy" section on page 38-16.

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**.

Step 2 If necessary, check the **Enable User Identity** check box to enable user identity.

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 includes 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 the "Configuring Active Directory Server Groups" section on page 38-12.

Step 6 Click **OK** to save the domain settings and close this dialog box.

What to Do Next

See the "Configuring Active Directory Server Groups" section on page 38-12 and the "Configuring Active Directory Agent Groups" section on page 38-13.

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 box appears.

Step 2 To add an Active Directory server group for the Identify Firewall, click **Add**.

The Add Active Directory Server Group dialog box appears.

Step 3 To add servers to an Active Directory server group, select the group from the Active Directory Server Groups list, then click **Add**.

The Add Active Directory Server dialog box appears.

Step 4 Click **OK** to save the settings and close this dialog box.

What to Do Next

See the "Configuring Active Directory Agents" section on page 38-12 and the "Configuring Active Directory Agent Groups" section on page 38-13.

Configuring Active Directory Agents

Prerequisites

Make sure that you have the following information before configuring the AD Agents:

- AD agent IP address
- Shared secret between the ASA and AD agent

To configure the AD Agents, perform the following steps:

Step 1 Choose **Configuration > Firewall > Identity Options**.

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. Alternatively, choose an agent group from the list and click **Edit**.

To continue, see the "Configuring Active Directory Agent Groups" section on page 38-13.

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

What to Do Next

Configure AD Agent groups. See the "Configuring Active Directory Agent Groups" section on page 38-13.

Configure access rules for the Identity Firewall. See the "Configuring Identity-Based Security Policy" section on page 38-16.

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 the 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 the ASA and AD Agent.

To configure the AD Agent Groups, perform the following steps:

Step 1 From the Configure Active Directory Agents dialog box, 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 the 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 perform to contact the AD Agent when it is not responding.

Step 8 Enter the shared secret key that is used between the secondary AD Agent and the ASA.

Step 9 Click **OK** to save your changes and close this dialog box.

What to Do Next

Configure access rules for the Identity Firewall. See the "Configuring Identity-Based Security Policy" section on page 38-16.

Configuring Identity Options

Use this pane to add or edit the Identity Firewall feature; check 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 meet the prerequisites for the AD Agent and Microsoft Active Directory. See the "Prerequisites" section on page 38-9 for the requirements of 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**.

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** to display the Add Domain dialog box.

Step 4 To continue, see the "Configuring the Active Directory Domain" section on page 38-11.

Step 5 For domains that have 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 is marked as disabled in the Monitoring > Properties > Identity > Users pane.

Step 6 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 that 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 incorrectly associates the user-IP mapping with the domain name that you entered 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 7 In the Active Directory Agent section, select the AD Agent group from the drop-down list. To add AD Agent groups, click **Manage**. For more information, see the "Configuring Active Directory Agents" section on page 38-12.

Step 8 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 9 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 10 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 because only users of received packets are queried and stored.

Step 11 In the Error Conditions section, select whether to disable rules if 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 the Monitoring > Properties > Identity > Users pane.

Step 12 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 13 In the Error Conditions section, choose whether or not 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 14 In the Error Conditions section, choose whether to track users that are not found.

Step 15 In the Users section, choose 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. After 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 16 In the NetBIOS Logout Probe section, enable NetBIOS probing and set the probe timer (from 1 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 17 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 includes the username 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 username 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 18 Click **Apply** to save the Identity Firewall configuration.

What to Do Next

Configure the Active Directory domain and server groups. See the "Configuring the Active Directory Domain" section on page 38-11 and the "Configuring Active Directory Server Groups" section on page 38-12.

Configure AD Agents. See the "Configuring Active Directory Server Groups" section on page 38-12.

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 38-8) can take advantage of an identity firewall. You can now add user identity arguments to extended ACLs, as well as network-based parameters.

Features that can use identity include the following:

- Access rules—An access rule permits or denies traffic on an interface using network information. With an identity firewall, you can control access based on user identity. See Chapter 6, "Configuring Access Rules," in the firewall configuration guide.
- AAA rules—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. 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 AAA 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. For example:

```
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
access-list 100 ex deny any any
access-group 100 in interface inside

access-list 200 ex deny ip user ANY any any
access-list 200 ex permit user NONE any any
aaa authenticate match 200 inside user-identity
```

  For more information, see Chapter 7, "Configuring AAA Rules for Network Access," in the firewall configuration guide.
- 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 25, "Configuring the ASA for Cisco Cloud Web Security," in the firewall configuration guide.

- VPN filter—Although a VPN does not support identity firewall ACLs in general, you can 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 (with the **no sysopt connection permit-vpn** command). You can also use an identity firewall ACL with the VPN filter feature; a VPN filter accomplishes a similar effect as allowing access rules in general.

Monitoring the Identity Firewall

This section includes the following topic:

- Monitoring AD Agents, page 38-17
- Monitoring Groups, page 38-17
- Monitoring Memory Usage for the Identity Firewall, page 38-18
- Monitoring Users for the Identity Firewall, page 38-18

Monitoring AD Agents

To monitor the AD Agent component of the Identity Firewall, perform the following steps:

Step 1 Choose **Monitoring > Properties > Identity > AD Agent**.

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

To monitor the user groups configured for the Identity Firewall, perform the following steps:

Step 1 Choose **Monitoring > Properties > Identity > Group**.

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 the list of user groups in the *domain\group_name* format.

Monitoring Memory Usage for the Identity Firewall

To monitor the memory usage that the Identity Firewall consumes on the ASA, perform the following steps:

Step 1 Choose **Monitoring > Properties > Identity > Memory Usage**.

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 login security logs.
- AD Agent
- Miscellaneous
- Total Memory Usage

Note How you configure the Identity Firewall to retrieve user information from the AD Agent affects the amount of memory used by the feature. You specify whether the ASA uses on-demand retrieval or full download retrieval. Choosing on-demand retrieval has the benefit of using less memory because only users of received packets are queried and stored. For more information, see the "Configuring Identity Options" section on page 38-13.

Monitoring Users for the Identity Firewall

To display information about all users contained in the IP-user mapping database used by the Identity Firewall, perform the following steps:

Step 1 Choose **Monitoring > Properties > Identity > User**.

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 the default domain is not specified, the default domain is LOCAL.

The idle time is stored on a per-user basis instead of by 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-in users have the disabled status. 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 tab lets you view important information about the traffic passing through your ASA. Choose the **Home > Firewall Dashboard > Top Usage Statistics > Top 10 Users** tab.

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—the Microsoft Active Directory and Cisco Active Directory (AD) Agent. For more information, see the "Configuring the Identity Firewall" section on page 38-10.

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 38-1 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 38-1 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 39-1
- Licensing Requirements for Cisco TrustSec, page 39-11
- Prerequisites for Using Cisco TrustSec, page 39-11
- Guidelines and Limitations, page 39-12
- Configuring the ASA for Cisco TrustSec Integration, page 39-14
- Monitoring Cisco TrustSec, page 39-21
- Feature History for the Cisco TrustSec Integration, page 39-22

Information About the ASA Integrated with Cisco TrustSec

This section includes the following topics:

- Information about Cisco TrustSec, page 39-2
- About SGT and SXP Support in Cisco TrustSec, page 39-2
- Roles in the Cisco TrustSec Feature, page 39-3
- Security Group Policy Enforcement, page 39-4
- How the ASA Enforces Security Group-Based Policies, page 39-4
- Effects of Changes to Security Groups on the ISE, page 39-6
- About Speaker and Listener Roles on the ASA, page 39-6
- SXP Chattiness, page 39-7
- SXP Timers, page 39-8
- IP-SGT Manager Database, page 39-8
- Features of the ASA-Cisco TrustSec Integration, page 39-9

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 employ 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 6-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 across the customers' networks, at the access, distribution, and core layers of the network, and in the data center.

Cisco TrustSec provides access control 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 feature, enforcement devices use 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 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 more information about Cisco TrustSec, see http://www.cisco.com/go/trustsec.

About SGT and SXP Support in Cisco TrustSec

In the Cisco TrustSec feature, security group access transforms a topology-aware network into a role-based network, which enables 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 ACL.

An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB), which occurs 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 ACLs. SXP, a control plane protocol, passes IP-SGT mapping 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 to initiate a connection. Additionally, an SXP connection is uniquely identified by the source and destination IP addresses.

Roles in the Cisco TrustSec Feature

To provide identity and policy-based access enforcement, the Cisco TrustSec feature includes the following roles:

- 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 feature, 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, Sensor IPS, and Communication Manager.

- Policy Administration Point (PAP)—A policy administration point defines and inserts policies into the 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 feature, 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 mapping 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 PEP role in the identity architecture. Using SXP, the ASA learns identity information directly from authentication points and uses it to enforce identity-based policies.

Security Group Policy Enforcement

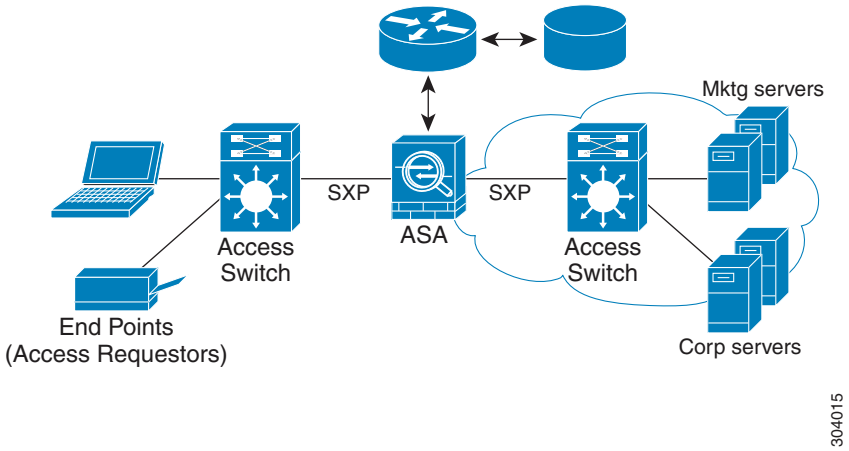
Security policy enforcement is based on security group name. An end-point device attempts to access a resource in the data center. Compared to traditional IP-based policies configured on firewalls, identity-based policies are configured based on user and device identities. For example, mktg-contractor 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 are 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 39-1 show a deployment for security group name-based policy enforcement.

Figure 39-1 Security Group Name-Based Policy Enforcement Deployment



Implementing Cisco TrustSec allows you to configure security policies that support server segmentation and includes the following features:

- 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 use 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 38, "Configuring the Identity Firewall" for information about configuring user-based security policies.

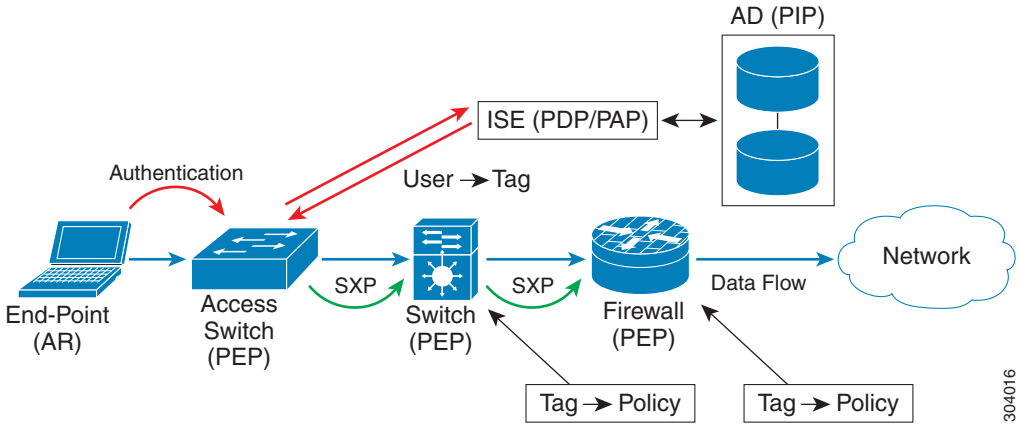
To configure the ASA to function with Cisco TrustSec, you must import a Protected Access Credential (PAC) file from the ISE. For more information, see the "Importing a Protected Access Credential (PAC) File" section on page 39-16.

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 (that is, 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 that the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names included in security policies that have been configured on it; then the ASA activates those security policies locally. If the ASA cannot resolve a security group name, it generates a syslog message for the unknown security group name.

Figure 39-2 shows how a security policy is enforced in Cisco TrustSec.

Figure 39-2 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 information 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 and looks up the SGTs for the source and destination IP addresses using the IP-SGT mapping passed by SXP.

   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 plane, tracks IP-SGT mapping 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 mapped entry.

   If the ASA is configured as a Speaker, the ASA transmits all IP-SGT mapping entries to its SXP peers. For more information, see the "About Speaker and Listener Roles on the ASA" section on page 39-6.
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 ASA that include 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 to be unknown and generates a syslog message. After the ASA refreshes the security group table from the ISE and learns the security group name, the ASA generates a syslog message indicating that the security group name is known.

Effects of Changes to Security Groups on the ISE

The ASA periodically refreshes the security group table by downloading an updated table from the ISE. Security groups can change on the ISE between downloads. These changes are not reflected on the ASA until it refreshes the security group table.

Tip We recommend that you schedule policy configuration changes on the ISE during a maintenance window, then manually refresh the security group table on the ASA to make sure the security group changes have been incorporated.

Handling policy configuration changes in this way maximizes the chances of security group name resolution and immediate activation of security policies.

The security group table is automatically refreshed when the environment data timer expires. You can also trigger a security group table refresh on demand.

If a security group changes on the ISE, the following events occur when the ASA refreshes the security group table:

- Only security group policies that have been configured using security group names need to be resolved with the security group table.Policies that include security group tags are always active.
- When the security group table is available for the first time, all policies with security group names are walked through, security group names are resolved, and policies are activated. All policies with tags are walked through, and syslogs are generated for unknown tags.
- If the security group table has expired, policies continue to be enforced according to the most recently downloaded security group table until you clear it, or a new table becomes available.
- When a resolved security group name becomes unknown on the ASA, it deactivates the security policy; however, the security policy persists in the ASA running configuration.
- If an existing secuity group is deleted on the PAP, a previously known security group tag can become unknown, but no change in policy status occurs on the ASA. A previously known security group name can become unresolved, and the policy is then inactivated. If the security group name is reused, the policy is recompiled using the new tag.
- If a new security group is added on the PAP, a previously unknown security group tag can become known, a syslog message is generated, but no change in policy status occurs. A previously unknown security group name can become resolved, and associated policies are then activated.
- If a tag has been renamed on the PAP, policies that were configured using tags display the new name, and no change in policy status occurs. Policies that were configured with security group names are recompiled using the new tag value.

About Speaker and Listener Roles on the ASA

The ASA supports SXP to send and receive IP-SGT mapping entries to and from other network devices. Using 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 mapping

entries from upstream devices (such as datacenter devices) back to 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 mapping entries collected on the ASA to upstream devices for policy enforcement.
- Listener mode—Configures the ASA so that it can receive IP-SGT mapping entries from downstream devices (SGT-capable switches) and use that information to create policy definitions.

If one end of an SXP connection is configured as a 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 fails and the ASA generates a syslog message.

Multiple SXP connections can learn IP-SGT mapping entries that have been downloaded from the IP-SGT mapping database. After an SXP connection to an SXP peer is established on the ASA, the Listener downloads the entire IP-SGT mapping database from the Speaker. All changes that occur after this are sent only when a new device appears on the network. As a result, the rate of SXP information flow is proportional to the rate at which end hosts authenticate to the network.

IP-SGT mapping entries that have been learned through SXP connections are maintained in the SXP IP-SGT mapping database. The same mapping entries may be learned through different SXP connections. The mapping database maintains one copy for each mapping entry learned. Multiple mapping entries of the same IP-SGT mapping value are identified by the peer IP address of the connection from which the mapping was learned. SXP requests that the IP-SGT Manager add a mapping entry when a new mapping is learned the first time and remove a mapping entry when the last copy in the SXP database is removed.

Whenever an SXP connection is configured as a Speaker, SXP requests that the IP-SGT Manager forward all the mapping entries collected on the device to the peer. When a new mapping is learned locally, the IP-SGT Manager requests that SXP forward it through connections that are configured as Speakers.

Configuring the ASA to be both a Speaker and a Listener for an SXP connection can cause SXP looping, which means that SXP data can be received by an SXP peer that originally transmitted it.

SXP Chattiness

The rate of SXP information flow is proportional to the rate at which end hosts authenticate into the network. After an SXP peering is established, the listener device downloads the entire IP-SGT database from the speaker device. After that, all changes are sent incrementally only when a new device appears on the network or leaves the network. Also, note that only access devices that are attached to the new device initiate this incremental update to the upstream device.

In other words, SXP protocol is no chattier than the authentication rate, which is limited to the capability of the authentication server. Therefore, SXP chattiness is not a major concern.

SXP Timers

- Retry Open Timer—The retry open timer is triggered if one SXP connection on the device is not up. After the retry open timer expires, the device goes through the entire connection database and if any connection is in the off or "pending on" state, the retry open timer restarts. The default timer value is 120 seconds. A zero value means the retry timer does not start. The retry open timer continues until all the SXP connections are set up, or the retry open timer has been configured to be 0.
- Delete Hold-Down Timer—The connection-specific delete hold-down timer is triggered when a connection on the Listener is torn down. The mapping entries that have been learned are not deleted immediately, but are held until the delete hold-down timer expires. The mapping entries are deleted after this timer expires. The delete hold-down timer value is set to 120 seconds and is not configurable.
- Reconciliation Timer—If an SXP connection is brought up within the delete hold-down timer period, a bulk update is performed on this connection. This means that the most recent mapping entries are learned and are associated with a new connection instantiation identifier. A periodic, connection-specific reconciliation timer starts in the background. When this reconciliation timer expires, it scans the entire SXP mapping database and identifies all mapping entries that have not been learned in the current connection session (that is, mapping entries with an unmatched connection instantiation identifier), and marks them for deletion. These entries are deleted in the subsequent reconciliation review. The default reconciliation timer value is 120 seconds. A zero value is not allowed on the ASA to prevent obsolete entries from staying for an unspecified length of time and causing unexpected results in policy enforcement.
- HA Reconciliation Timer—When HA is enabled, the SXP mapping database of the active and standby units are in sync. The new active unit tries to establish new SXP connections to all its peers and acquires the latest mapping entries. An HA reconciliation timer provides a way of identifying and removing old mapping entries. It starts after a failover occurs, which gives the ASA time to acquire the latest mapping entries. After the HA reconciliation timer expires, the ASA scans the entire SXP mapping database and identifies all the mapping entries have not been learned in the current connection session. Mapping entries with unmatched instantiation identifiers are marked for deletion. This reconciliation mechanism is the same as that of the reconciliation timer. The time value is the same as the reconciliation timer and is configurable.

After an SXP peer terminates its SXP connection, the ASA starts a delete hold-down timer. Only SXP peers designated as Listeners can terminate a connection. If an SXP peer connects while the delete hold-down timer is running, the ASA starts the reconciliation timer; then the ASA updates the IP-SGT mapping database to learn the most recent mapping.

IP-SGT Manager Database

The IP-SGT Manager database does not synchronize any entries from the active unit to the standby unit. Each source from which the IP-SGT Manager database receives IP-SGT mapping entries synchronizes its database from the active unit to the standby unit, then provides the final IP-SGT mapping to the IP-SGT Manager on the standby unit.

For Version 9.0(1), the IP-SGT Manager database receives IP-SGT mapping updates from the SXP source only.

Features of the ASA-Cisco TrustSec Integration

The ASA includes Cisco TrustSec as part of its identity-based firewall feature. Cisco TrustSec provides the following capabilities:

Flexibility

- The ASA can be configured as an SXP Speaker or Listener, or both.

  See the "About Speaker and Listener Roles on the ASA" section on page 39-6.
- The ASA supports SXP for IPv6 and IPv6-capable network devices.
- SXP can change mapping entries for IPv4 and IPv6 addresses.
- SXP end points support IPv4 and IPv6 addresses.
- The ASA supports SXP Version 2 only.
- 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 both the Active/Active and Active/Standby configurations.
- The ASA can communicate with the ISE configured for high availability (HA).
- You can configure multiple ISE servers on the ASA and if the first server is unreachable, it continues to the next server, and so on. However, if the server list is downloaded as part of the Cisco TrsutSec environment data, it is ignored.
- If the PAC file downloaded from the ISE expires on the ASA and it 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.

Clustering

- For Layer 2 networks, all units share the same IP address. When you change the interface address, the changed configuration is sent to all other units. When the IP address is updated from the interface of a particular unit, a notification is sent to update the IP-SGT local database on this unit.
- For Layer 3 networks, a pool of addresses is configured for each interface on the master unit, and this configuration is synchronized to the slave units. On the master unit, a notification of the IP addresses that have been assigned to the interface is sent, and the IP-SGT local database is updated. The IP-SGT local database on each slave unit can be updated with the IP address information for the master unit by using the address pool configuration that has been synchronized to it, where the first address in the pool for each interface always belongs to the master unit.

  When a slave unit boots, it notifies the master unit. Then the master unit goes through the address pool on each interface and computes the IP address for the new slave unit that sent it the notification, and updates the IP-SGT local database on the master unit. The master unit also notifies the other slave units about the new slave unit. As part of this notification processing, each slave unit computes

the IP address for the new slave unit and adds this entry to the IP-SGT local database on each slave unit. All the slave units have the address pool configuration to determine the IP address value. For each interface, the value is determined as follows:

Master IP + (M-N), where:

M—Maximum number of units (up to 8 are allowed)

N—Slave unit number that sent the notification

When the IP address pool changes on any interface, the IP addresses for all the slave units and the master unit need to be recalculated and updated in the IP-SGT local database on the master unit, as well as on every other slave unit. The old IP address needs to be deleted, and the new IP address needs to be added.

When this changed address pool configuration is synchronized to the slave unit, as a part of configuration change processing, each slave unit recomputes the IP address for the master unit and for every other slave unit whose IP address has changed, then removes the entry for the old IP address and adds the new IP address.

Scalability

Table 39-1 show the number of IP-SGT mapping entries that the ASA supports.

Table 39-1 Capacity Numbers for IP-SGT Mapping Entries

ASA Model	Number of IP-SGT Mapping 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

Table 39-2 shows the number of SXP connections that the ASA supports.

Table 39-2 SXP Connections

ASA Model	Number of SXP TCP Connections
5505	10
5510	25
5520	50
5540	100
5550	150
5580-20	250

Table 39-2 SXP Connections (continued)

ASA Model	Number of SXP TCP Connections
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 for Cisco TrustSec

Model	License Requirement
All models	Base License.

Prerequisites for Using Cisco TrustSec

Before configuring the ASA to use Cisco TrustSec, you must perform the following tasks:

- Registering the ASA with the ISE, page 39-11
- Creating a Security Group on the ISE, page 39-12
- Generating the PAC File, page 39-12

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. To register the ASA with the ISE, perform the following steps:

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, enter a shared secret in the Authentication Settings area.

   When you configure the AAA sever on the ASA, provide the shared secret that 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 how 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. To create a security group on the ISE, perform the following steps:

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 device ID credentials and a password for the ASA.

Generating the PAC File

Before generating the PAC file, you must have registered the ASA with the ISE. To generate the PAC file, perform the following steps:

1. Log into the ISE.
2. Choose **Administration > Network Resources > Network Devices**.
3. From the list of devices, choose the ASA.
4. Under the Security Group Access (SGA), click **Generate PAC**.
5. To encrypt the PAC file, enter a password.

   The password (or encryption key) that 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 file from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC file does not have to reside on the ASA flash before you can import it.)

For information about the PAC file, see the "Importing a Protected Access Credential (PAC) File" section on page 39-16.

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 SXP endpoints.

Clustering Guidelines

Supported on the master unit and on slave units in a clustering environment.

Failover Guidelines

Supports a list of servers via configuration. If the first server is unreachable, the ASA tries 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.

Supports both Active/Standby and Active/Active scenarios. All SXP data is replicated from the active unit to the standby unit after it takes over.

Additional Guidelines

Cisco TrustSec supports the Smart Call Home feature in single context and multi-context mode, but not in the system context.

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 mapping on the device.
- NAT is not supported in SXP messages.
- SXP conveys IP-SGT mapping 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 that it uploads is invalid, and an IP-SGT mapping database lookup on the enforcement device does not yield valid results. As a result, the ASA cannot apply security group-aware security policy on the enforcement device.
- You can configure a default password for the ASA to use for SXP connections, or you can choose not to use a password; however, connection-specific passwords are not supported for SXP peers. The configured default SXP password should be consistent across the deployment network. If you configure a connection-specific password, connections may fail and a warning message appears. If you configure the connection with the default password, but it is not configured, the result is the same as when you have configured the connection with no password.
- 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-SGT mapping 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. In addition, if 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 into the ASA. For information about the PAC file, see the "Generating the PAC File" section on page 39-12 and the "Importing a Protected Access Credential (PAC) File" section on page 39-16.
- PAC files have expiration dates. You must import the updated PAC file before the current PAC file expires; otherwise, the ASA cannot 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 syslog message to indicate that those security policies changed.

See the "Refreshing Environment Data" section on page 39-20 for information about manually updating the security group table on the ASA to include 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 to configure SXP connections. Create a TCP state bypass policy for traffic destined to SXP port TCP 64999 among the SXP peers. Then apply the policy on the appropriate interfaces.

  For example, the following set of commands shows how to configure the ASA 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 includes the following topics:

- Task Flow for Configuring the ASA to Integrate with Cisco TrustSec, page 39-15
- Configuring the AAA Server for Cisco TrustSec Integration, page 39-15
- Importing a Protected Access Credential (PAC) File, page 39-16
- Configuring the Security Exchange Protocol (SXP), page 39-17
- Adding an SXP Connection Peer, page 39-19
- Refreshing Environment Data, page 39-20
- Configuring the Security Policy, page 39-20

Task Flow for Configuring the ASA to Integrate with Cisco TrustSec

Prerequisite

Before configuring the ASA to integrate with Cisco TrustSec, you must complete the following tasks:

- Register the ASA with the ISE.
- Create a security group on the ISE.
- Generate the PAC file on the ISE to import into the ASA.

See the "Prerequisites for Using Cisco TrustSec" section on page 39-11 for more information.

To configure the ASA to integrate with Cisco TrustSec, perform the following tasks:

Step 1 Configure the AAA server.

See the "Configuring the AAA Server for Cisco TrustSec Integration" section on page 39-15.

Step 2 Import the PAC file from the ISE.

See the "Importing a Protected Access Credential (PAC) File" section on page 39-16.

Step 3 Enable and set the default values for SXP.

See the "Configuring the Security Exchange Protocol (SXP)" section on page 39-17.

Step 4 Add SXP connection peers for the Cisco TrustSec architecture.

See the "Adding an SXP Connection Peer" section on page 39-19.

Step 5 As necessary, refresh environment data for the ASA.

See the "Refreshing Environment Data" section on page 39-20.

Step 6 Configure the security policy.

See the "Configuring the Security Policy" section on page 39-20.

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.

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 configuration fails.
- 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 to obtain 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**.

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 cannot communicate with the ISE. Contact your ISE administrator to obtain 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 the "Configuring RADIUS Server Groups" section on page 34-15.

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 the "Adding a RADIUS Server to a Group" section on page 34-16.

Step 9 Click **OK**. The ASA adds the ISE server to the list of AAA servers.

Step 10 Click **Apply** 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 (that is, 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 before the RADIUS transaction. The ASA initiates a RADIUS transaction with the ISE using the PAC file for authentication.

Tip The PAC file includes 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 downloads 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 only works on the ASA when the file was generated by a correctly configured ISE. See the "Registering the ASA with the ISE" section on page 39-11 for more information.
- 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 file from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC file does not need 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 file 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**.

Step 2 Check the **Enable Security Exchange Protocol** check box 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, reenter the password to confirm it.

Step 7 Click **Import**.

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

Note For SXP to be operational on the ASA, at least one interface must be in the UP/UP state.

Currently, when SXP is enabled with all interfaces down, the ASA does not display a message indicating that SXP is not working or it could not be enabled. If you check the configuration by entering the **show running-config** command, the command output displays the following message:

```
"WARNING: SXP configuration in process, please wait for a few moments and try again."
```

This message is generic and does not specify the reason why SXP is not working.

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**.

Step 2 Check the **Enable Security Exchange Protocol** check box to enable SXP. By default, SXP is disabled.

In multi-context mode, you enable SXP 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 syslog 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 continues to try 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 in the range of 0 to 64000 seconds. If you specify 0 seconds, the timer never expires and the ASA does not try to connect to SXP peers. By default, the timer value 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 value.

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 mapping.

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 in the range of 1 to 64000 seconds. By default, the timer value is 120 seconds.

Note You cannot specify 0 seconds for the timer, because this value would prevent the reconcile timer from starting. Not allowing the reconcile timer to run would keep stale entries for an undefined period and cause unexpected results from 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**.

Step 2 If necessary, check the **Enable Security Exchange Protocol** check box 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 choosing one of the following values:

- Default—Use the default password configured for SXP connections.

  See the "Configuring the Security Exchange Protocol (SXP)" section on page 39-17.
- 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 choosing 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 mapping to upstream devices.
- Listener—The ASA can receive IP-SGT mapping from downstream devices.

See the "About Speaker and Listener Roles on the ASA" section on page 39-6.

Step 9 Click **OK**. The peer appears in the Connection Peers list.

Step 10 Click **Apply** to save your 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 that is 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 do 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, so refresh the data on the ASA to make sure that any security group changes 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 the environment data, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Firewall > Identity By TrustSec**.

Step 2 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 access rule, Chapter 6, “Configuring Access Rules.”

- To configure security group object groups, which can be used in the ACL, see the "Configuring Local User Groups" section on page 20-7.

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 6, "Configuring Access Rules," in the firewall configuration guide. 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, and so on), 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 Cisco TrustSec

To monitor Cisco TrustSec on the ASA, choose one of the following paths in ASDM:

Path	Purpose
Monitoring > Properties > Identity By TrustSec > SXP Connections	Displays the configured default values for the Cisco TrustSec infrastructure and the SXP commands.
Monitoring > Properties > Connections	Displays data for all SXP connections. Filters the IP address-security group table mapping entries so that you view the data by security group table value, security group name, or IP address.
Monitoring > Properties > Identity By TrustSec > Environment Data	Displays the Cisco TrustSec environment information contained in the security group table on the ASA.
Monitoring > Properties > Identity By TrustSec > IP Mapping	Displays the IP address-security group table mapping entries from the IP address-security group table mapping database maintained in the datapath. Filters the IP address-security group table mapping entries so that you view the data by security group table value, security group name, or IP address. **Tip** Click **Where Used** to display where the selected security group object is used in an ACL or nested in another security group object.
Monitoring > Properties > Identity By TrustSec > PAC	Displays information about the PAC file imported into the ASA from the ISE. Displays a warning message when the PAC file has expired or is within 30 days of expiring.

Feature History for the Cisco TrustSec Integration

Table 39-3 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 39-3 Feature History for the Cisco TrustSec Integration

Feature Name	Platform Releases	Feature Information
Cisco TrustSec Integration	9.0(1)	Cisco TrustSec provides access control 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 feature, enforcement devices use 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 use the Cisco TrustSec feature for other types of security group-based policies, such as application inspection; for example, you can configure a class map that includes 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 40-1
- Licensing Requirements for Digital Certificates, page 40-9
- Prerequisites for Local Certificates, page 40-10
- Guidelines and Limitations, page 40-10
- Configuring Digital Certificates, page 40-11
- Configuring CA Certificate Authentication, page 40-12
- Monitoring CRLs, page 40-20
- Configuring Identity Certificates Authentication, page 40-24
- Configuring Code Signer Certificates, page 40-29
- Authenticating Using the Local CA, page 40-31
- Managing the User Database, page 40-35
- Managing User Certificates, page 40-37
- Monitoring CRLs, page 40-38
- Feature History for Certificate Management, page 40-39

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 40-12.
- 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 40-24.
- *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 40-29.

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 40-31, the "Managing the User Database" section on page 40-35, and the "Managing User Certificates" section on page 40-37.

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 40-14.

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 40-3
- Certificate Scalability, page 40-3
- Key Pairs, page 40-4
- Trustpoints, page 40-4
- Revocation Checking, page 40-5

- The Local CA, page 40-7
- Using Certificates and User Login Credentials, page 40-8

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. We recommend using a key size of at least 2048.
- 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.

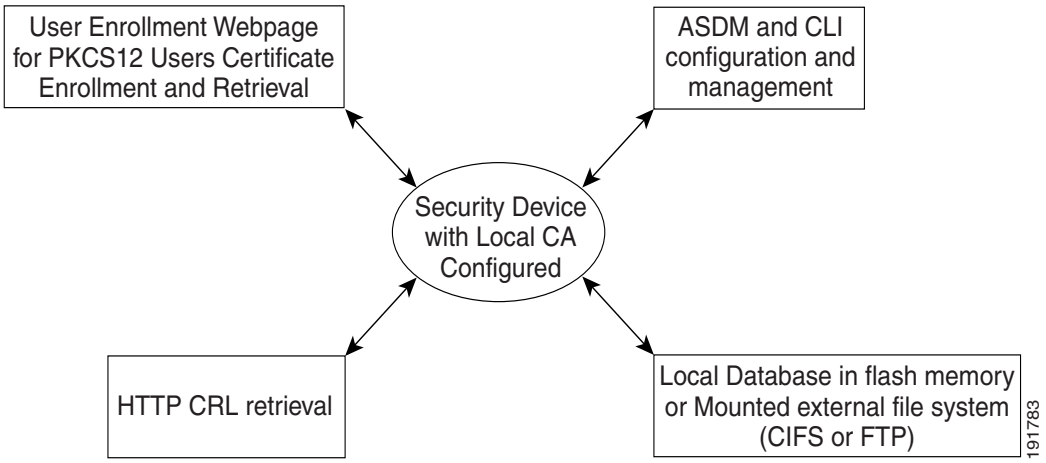
No limits exist on the number of users that can be stored in the local CA user database; however, if flash memory storage issues arise, syslogs are generated to alert the administrator to take action, and the local CA could be disabled until the storage issues are resolved. Flash memory can store a database 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 40-1, 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 40-1 The Local CA



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 40-8
- Using Certificates, page 40-9

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 Digital Certificates

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 16-1.
- 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 16-3.

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 failover for local CAs.

IPv6 Guidelines

Does not support 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. We recommend using a key size of at least 2048.
- 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.
- Use only valid characters and values for certificate parameters when you import them on the ASA.
- To use a wildcard (*) symbol, make sure that you use encoding on the CA server that allows this character in the string value. Although RFC 5280 recommends using either a UTF8String or PrintableString, you should use UTF8String because PrintableString does not recognize the wildcard as a valid character. The ASA rejects the imported certificate if an invalid character or value is found during the import. For example:

```
ERROR: Failed to parse or verify imported certificate ciscoasa(config)# Read
162*H÷ytes as CA certificate:0U0= \Ivr"phÕV°3é¼þ0 CRYPTO_PKI(make trustedCerts list)
CERT-C: E ../cert-c/source/certlist.c(302) : Error #711h
CRYPTO_PKI: Failed to verify the ID certificate using the CA certificate in trustpoint
mm.
CERT-C: E ../cert-c/source/p7contnt.c(169) : Error #703h
crypto_certc_pkcs7_extract_certs_and_crls failed (1795):
crypto_certc_pkcs7_extract_certs_and_crls failed
CRYPTO_PKI: status = 1795: failed to verify or insert the cert into storage
```

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 40-12

- Configuring CA Certificates for Revocation, page 40-20
- Configuring CRL Retrieval Policy, page 40-21
- Configuring CRL Retrieval Methods, page 40-21
- Configuring OCSP Rules, page 40-22
- Configuring Advanced CRL and OCSP Settings, page 40-23

This section describes how to configure digital certificates for the ASA Services Module and includes the following topics:

- Configuring CA Certificate Authentication, page 40-14
- Adding or Installing a CA Certificate, page 40-15
- Editing or Removing a CA Certificate Configuration, page 40-15
- Showing CA Certificate Details, page 40-16
- Configuring CA Certificates for Revocation, page 40-16
- Configuring CRL Retrieval Policy, page 40-17
- Configuring CRL Retrieval Methods, page 40-17
- Configuring OCSP Rules, page 40-18
- Configuring Advanced CRL and OCSP Settings, page 40-19

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 40-13
- Editing or Removing a CA Certificate Configuration, page 40-14
- Showing CA Certificate Details, page 40-14

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 40-14.

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 40-21
- "Configuring CRL Retrieval Methods" section on page 40-21
- "Configuring OCSP Rules" section on page 40-22
- "Configuring Advanced CRL and OCSP Settings" section on page 40-23

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 40-15.

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 40-17
- "Configuring CRL Retrieval Methods" section on page 40-17
- "Configuring OCSP Rules" section on page 40-18
- "Configuring Advanced CRL and OCSP Settings" section on page 40-19

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 40-17.

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 40-17.

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 40-18.

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 40-19.

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 40-20.

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 40-21.

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 40-21.

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 40-22.

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 40-23.

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 40-24.

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 40-24
- Showing Identity Certificate Details, page 40-26
- Deleting an Identity Certificate, page 40-27
- Exporting an Identity Certificate, page 40-27
- Generating a Certificate Signing Request, page 40-27
- Installing Identity Certificates, page 40-28

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 40-24. 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 40-24.

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 40-29.

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 40-30
- Deleting a Code Signer Certificate, page 40-30
- Importing a Code Signer Certificate, page 40-30
- Exporting a Code Signer Certificate, page 40-31

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 40-31.

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 40-32
- Deleting the Local CA Server, page 40-34

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. We recommend that you use a key size of at least 2048.

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. We recommend that you use a key size of at least 2048.

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 40-35.

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 40-35
- Sending an Initial OTP or Replacing OTPs, page 40-36
- Editing a Local CA User, page 40-36
- Deleting a Local CA User, page 40-37
- Allowing User Enrollment, page 40-37
- Viewing or Regenerating an OTP, page 40-37

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 40-37.

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 40-38.

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 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. 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 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 8

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 41-1
- Licensing Requirements for Logging, page 41-5
- Prerequisites for Logging, page 41-6
- Guidelines and Limitations, page 41-6
- Configuring Logging, page 41-7
- Monitoring the Logs, page 41-25
- Feature History for Logging, page 41-28

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 41-2
- Analyzing Syslog Messages, page 41-2
- Syslog Message Format, page 41-3
- Severity Levels, page 41-3
- Message Classes and Range of Syslog IDs, page 41-4
- Filtering Syslog Messages, page 41-4
- Sorting in the Log Viewers, page 41-4
- Using Custom Message Lists, page 41-5
- Using Clustering, page 41-5

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 41-1 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 41-1 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 41-1 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.
- The syslog server should be reachable through the ASA and ASASM. You should configure the ASA SM to deny ICMP unreachable messages on the interface through which the syslog server is reachable and to send syslogs to the same server. Make sure that you have enabled logging for all severity levels. To prevent the syslog server from crashing, suppress the generation of syslogs 313001, 313004, and 313005.
- When you use a custom message list to match only access list hits, the access list logs are not generated for access lists that have had their logging severity level increased to debugging (level 7). The default logging severity level is set to 6 for the **logging list** command. This default behavior is by design. When you explicitly change the logging severity level of the access list configuration to debugging, you must also change the logging configuration itself.

  The following is sample output from the **show running-config logging** command that will not include access list hits, because their logging severity level has been changed to debugging:

```
ciscoasa# show running-config logging
logging enable
logging timestamp
logging list test message 106100
logging buffered test
```

The following is sample output from the **show running-config logging** command that will include access list hits:

```
ciscoasa# show running-config logging
logging enable
logging timestamp
logging buffered debugging
```

In this case, the access list configuration does not change and the number of access list hits appears, as shown in the following example:

```
ciscoasa(config)# access-list global line 1 extended permit icmp any host 4.2.2.2 log
debugging interval 1 (hitcnt=7) 0xf36b5386
ciscoasa(config)# access-list global line 2 extended permit tcp host 10.1.1.2 any eq
www log informational interval 1 (hitcnt=18) 0xe7e7c3b8
ciscoasa(config)# access-list global line 3 extended permit ip any any (hitcnt=543)
0x25f9e609
```

Configuring Logging

This section describes how to configure logging and includes the following topics:

- Enabling Logging, page 41-7
- Configuring an Output Destination, page 41-8

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:

Step 1 In ASDM, 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**

Step 2 Check the **Enable logging** check box to turn on logging.

What to Do Next

See the "Configuring an Output Destination" section on page 41-8.

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 41-9
- Configuring FTP Settings, page 41-10
- Configuring Logging Flash Usage, page 41-10
- Configuring Syslog Messaging, page 41-10
- Editing Syslog ID Settings, page 41-11
- Including a Device ID in Non-EMBLEM Formatted Syslog Messages, page 41-12
- Sending Syslog Messages to the Internal Log Buffer, page 41-12
- Saving an Internal Log Buffer to Flash, page 41-13
- Viewing and Copying Logged Entries with the ASDM Java Console, page 41-13
- Sending Syslog Messages to an E-mail Address, page 41-14
- Adding or Editing E-Mail Recipients, page 41-14
- Configuring the Remote SMTP Server, page 41-15
- Viewing Syslog Messages in ASDM, page 41-15
- Applying Message Filters to a Logging Destination, page 41-15
- Applying Logging Filters, page 41-16
- Adding or Editing a Message Class and Severity Filter, page 41-17
- Adding or Editing a Syslog Message ID Filter, page 41-17
- Sending Syslog Messages to the Console Port, page 41-17
- Sending Syslog Messages to a Telnet or SSH Session, page 41-18
- Creating a Custom Event List, page 41-18
- Generating Syslog Messages in EMBLEM Format to a Syslog Server, page 41-19
- Adding or Editing Syslog Server Settings, page 41-19
- Generating Syslog Messages in EMBLEM Format to Other Output Destinations, page 41-20
- Changing the Amount of Internal Flash Memory Available for Logs, page 41-20
- Configuring the Logging Queue, page 41-21
- Sending All Syslog Messages in a Class to a Specified Output Destination, page 41-21
- Enabling Secure Logging, page 41-21
- Including the Device ID in Non-EMBLEM Format Syslog Messages, page 41-22

- Including the Date and Time in Syslog Messages, page 41-22
- Disabling a Syslog Message, page 41-22
- Changing the Severity Level of a Syslog Message, page 41-23
- Limiting the Rate of Syslog Message Generation, page 41-23
- Assigning or Changing Rate Limits for Individual Syslog Messages, page 41-24
- Adding or Editing the Rate Limit for a Syslog Message, page 41-24
- Editing the Rate Limit for a Syslog Severity Level, page 41-24

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 41-10.

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 41-10.

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 41-11 and the "Including a Device ID in Non-EMBLEM Formatted Syslog Messages" section on page 41-12.

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

You need to specify 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 send syslog messages to the internal log buffer, perform the following steps:

Step 1 To specify which syslog messages should be sent to the internal log buffer, choose one of the following:

- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
- **Configuration > Device Management > Logging > Logging Filters**

Step 2 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 3 To change the size of the internal log buffer, choose **Configuration > Device Management > Logging > Logging Setup**. The default buffer size is 4 KB.

The ASA and ASASM continue to save new messages to the internal log buffer and save the full log buffer content to internal flash memory. 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.

Step 4 To save new messages to another location, choose one of the following options:

- To send new messages to internal flash memory, check the **Flash** check box, then click **Configure Flash Usage**. The Configure Logging Flash Usage dialog box appears.
  a. Specify the maximum amount of flash memory in KB that you want to use for logging.
  b. Specify the minimum amount of free space in KB that logging will preserve in flash memory.
  c. Click **OK** to close this dialog box.
- To send new messages to an FTP server, check the **FTP Server** check box, then click **Configure FTP Settings**. The Configure FTP Settings dialog box appears.
  a. Check the **Enable FTP Client** check box.
  b. Enter the following information in the fields provided: FTP server IP address, path, username, and password.
  c. Confirm the password, then click **OK** to close this dialog box.

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, 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 Per-Process CPU Usage section in the firewall configuration guide.

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 41-14.

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 41-16.

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 41-14.

Step 6 Click **OK** to close this dialog box.

Step 7 To continue, see the "Configuring the Remote SMTP Server" section on page 41-15.

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 41-17 and the "Adding or Editing a Syslog Message ID Filter" section on page 41-17.

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 41-16.

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 41-18.

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 41-11.

Sending Syslog Messages to the Console Port

To send syslog messages to the console port, perform the following steps:

Step 1 In ASDM, choose one of the following options:

- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
- **Configuration > Device Management > Logging > Logging Filters**

Step 2 Select the console in the Logging Destination column, then click **Edit**.

The Edit Logging Filters dialog box appears.

Step 3 To specify which syslog messages should be sent to the console port, choose either syslogs from all event classes or syslogs from specific event classes.

Step 4 To continue, see the "Applying Logging Filters" section on page 41-16.

Sending Syslog Messages to a Telnet or SSH Session

To send syslog messages to a Telnet or SSH session, perform the following steps:

Step 1 In ASDM, choose one of the following:

- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
- **Configuration > Device Management > Logging > Logging Filters**

Step 2 Select the Telnet and SSH Sessions in the Logging Destination column, then click **Edit**.

The Edit Logging Filters dialog box appears.

Step 3 To specify which syslog messages should be sent to a Telnet or an SSH session, choose either syslogs from all event classes or syslogs from specific event classes.

Step 4 To continue, see the "Applying Logging Filters" section on page 41-16.

Step 5 To enable logging for the current session only, choose **Configuration > Device Management > Logging > Logging Setup.**

Step 6 Check the **Enable logging** check box, then click **Apply**.

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.

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. If you specify TCP, the ASA or ASASM discovers when the syslog server fails and as a security protection, new connections through the ASA are blocked. If you specify UDP, the ASA or ASASM continues to allow new connections whether or not the syslog server is operational. Valid port values for either protocol are 1025 through 65535. The default UDP port is 514. The default TCP port is 1470.

Note Sending syslogs over TCP is not supported on a standby ASA.

Step 4 To continue, see the "Adding or Editing Syslog Server Settings" section on page 41-19.

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.

Choose the protocol (either TCP or UDP) that is used by the syslog server to communicate with the ASA or ASASM. You can configure the ASA and ASASM to send data to a syslog server using either UDP or TCP, but not both. The default protocol is UDP if you do not specify a protocol.

Step 3 Enter the port number used by the syslog server to communicate with the ASA or ASASM.

Step 4 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 5 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 6 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, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Logging Setup**.

Step 2 Check the **Send syslogs in EMBLEM format** check box.

Step 3 To continue, see the "Applying Logging Filters" section on page 41-16.

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:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Logging Setup**.

Step 2 Check the **Enable Logging** check box.

Step 3 In the Logging to Internal Buffer area, check the **Save Buffer to Flash** check box.

Step 4 Click **Configure Flash Usage**.

The Configure Logging Flash Usage dialog box appears.

Step 5 Enter the maximum amount of flash memory in KB allowed to be used for logging.

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 6 Enter the minimum amount of free space in KB to be preserved for logging in flash memory.

Step 7 Click **OK** to close this dialog box.

Configuring the Logging Queue

To configure the logging queue, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Logging Setup.**

Step 2 Check the **Enable logging** check box.

Step 3 In the ASDM Logging area, enter 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

Step 4 Click **OK** to close this dialog box.

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, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Logging Filters.**

Step 2 To overrides the configuration in the specified output destination, choose the output destination that you want to change, then click **Edit**.

The Edit Logging Filters dialog box appears.

Step 3 Revise the settings in either the Syslogs from All Event Classes or Syslogs from Specific Event Classes area, then click **OK** to close this dialog box.

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, select a different filtering option for each output destination.

Enabling Secure Logging

To enable secure logging, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Syslog Server.**

Step 2 Select a syslog server for which you want to enable secure logging, then click **Edit**.

The Edit Syslog Server dialog box appears.

Step 3 Click the TCP radio button.

Note Secure logging does not support UDP; an error occurs if you try to use this protocol.

Step 4 Check the **Enable secure syslog with SSL/TLS** check box, then click **OK**.

Including the Device ID in Non-EMBLEM Format Syslog Messages

To include the device ID in non-EMBLEM format syslog messages, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration.**

Step 2 Check the **Enable syslog device ID** check box.

Step 3 In the Device ID area, click the **Hostname**, **Interface IP Address**, or **String** radio button.

- If you chose the Interface IP Address option, make sure that the correct interface is selected in the drop-down list.
- If you chose the String option, enter the device ID in the User-Defined ID field. The string can include as many as 16 characters.

Note If enabled, the device ID does not appear in EMBLEM-formatted syslog messages nor in SNMP traps.

Step 4 Click **OK** to close the Advanced Syslog Configuration dialog box.

Including the Date and Time in Syslog Messages

To include the date and time in syslog messages, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Syslog Setup.**

Step 2 In the Syslog ID Setup area, check the **Include timestamp in syslogs** check box.

Step 3 Click **Apply** to save your changes.

Disabling a Syslog Message

To disable a specified syslog message, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Syslog Setup**.

Step 2 Select the syslog that you want to disable from the table, t hen click **Edit**.

The Edit Syslog ID Settings dialog box appears.

Step 3 Check the **Disable messages** check box, then click **OK**.

Changing the Severity Level of a Syslog Message

To change the severity level of a syslog message, perform the following steps:

Step 1 In ASDM, choose **Configuration > Device Management > Logging > Syslog Setup**.

Step 2 Select the syslog whose severity level you want to change from the table, t hen click **Edit**.

The Edit Syslog ID Settings dialog box appears.

Step 3 Choose the desired severity level from the Logging Level drop-down list, then click **OK**.

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
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 41-24.

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 41-24.

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 41-24.

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 41-25
- Editing Filtering Settings, page 41-27
- Executing Certain Commands Using the Log Viewers, page 41-27

To monitor the logs in the log buffer or in real-time and assist in monitoring the system performance, eperform the following steps:

Step 1 In ASDM, choose one of the following:

- **Monitoring > Logging > Log Buffer > View**
- **Monitoring > Logging > Real-Time Log Viewer > View**

The Real-Time Log Viewer or Log Buffer dialog box that appears displays the message explanations, additional details, and recommended actions to take, if necessary, to resolve an error in a separate window.

Step 2 To continue, see the "Filtering Syslog Messages Through the Log Viewers" section on page 41-25.

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 41-2 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 41-2 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 41-2 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 SNMP

This chapter describes how to configure SNMP to monitor the ASA and ASASM and includes the following sections:

- Information About SNMP, page 42-1
- Licensing Requirements for SNMP, page 42-4
- Prerequisites for SNMP, page 42-4
- Guidelines and Limitations, page 42-4
- Configuring SNMP, page 42-6
- Monitoring SNMP, page 42-11
- Where to Go Next, page 42-12
- Additional References, page 42-12
- Feature History for SNMP, page 42-15

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 42-2
- SNMP Version 3, page 42-2

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 42-1 lists the terms that are commonly used when working with SNMP:

Table 42-1 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 42-3
- Security Models, page 42-3
- SNMP Groups, page 42-3
- SNMP Users, page 42-3
- SNMP Hosts, page 42-3
- Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software, page 42-4

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.
- You can add up to 4000 hosts. However, only 128 of this number can be for traps.
- The total number of active polling destinations supported is 128.

- You can specify a network object to indicate the individual hosts that you want to add as a host group.
- You can associate more than one user with one host.
- You can specify overlapping network objects in different **host-group** commands. The values that you specify for the last host group take effect for the common set of hosts in the different network objects.
- If you delete a host group or hosts that overlap with other host groups, the hosts are set up again using the values that have been specified in the configured host groups.
- The values that the hosts acquire depend on the specified sequence that you use to run the commands.

Configuring SNMP

This section describes how to configure SNMP and includes the following topics:

- Enabling SNMP, page 42-6
- Configuring an SNMP Management Station, page 42-7
- Configuring SNMP Traps, page 42-7
- Using SNMP Version 1 or 2c, page 42-8
- Using SNMP Version 3, page 42-9

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 42-7.

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 42-7.

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 42-8.
- See the "Using SNMP Version 3" section on page 42-9.

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 42-11.

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, on the SNMPv3 User/Group tab, click **Add > SNMP User**. To change user parameters, click **Edit > SNMP User**. To remove a configured user from a group, select the user, then click **Delete > SNMP User**. 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.

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 42-11.

Configuring a Group of Users

To configure an SNMP user list with a group of specified users in it, 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 group or a new user group, on the SNMPv3 User/Group tab, click **Add > SNMP User Group**. To change group parameters, click **Edit > SNMP Group**. To remove a configured user group, select it, then click **Delete > SNMP Group**. When you remove the last user in a group, ASDM deletes the group.

The Add SNMP User Group dialog box appears.

Step 3 Enter the user group name.

Step 4 To select an existing user or user group, click the **Existing User/User Group** radio button.

Step 5 To create a new user, click the **Create new user** radio button.

Step 6 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 7 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 8 Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.

Step 9 Indicate the type of authentication you want to use by clicking one of the two radio buttons: **MD5** or **SHA**.

Step 10 In the Authentication Password field, type the password to use for authentication.

Step 11 Retype the password to use for authentication.

Step 12 Indicate the type of encryption you want to use by clicking one of these three radio buttons: **DES**, **3DES**, or **AES**.

Step 13 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 14 Retype the password to use for encryption.

Step 15 Click **Add** to add the new user to the specified user group in the Members in Group pane. Click **Remove** to remove the new user from the Members in Group pane.

Step 16 Click **OK** to create a new user for the specified user group.

The Add SNMP User Group dialog box closes.

Step 17 Click **Apply**.

SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.

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 42-11
- SNMP Monitoring, page 42-11

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 host-group** command, then click **Send**.	Shows SNMP host group configurations.
Tools > Command Line Interface Enter the **show running-config snmp-server user** command, then click **Send**.	Shows SNMP user-based configuration settings.

Path	Purpose
Tools > Command Line Interface Enter the **show running-config snmp-server user-list** command, then click **Send**.	Shows SNMP user list configurations.
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 41, "Configuring Logging."

Additional References

For additional information related to implementing SNMP, see the following sections:

- RFCs for SNMP Version 3, page 42-12
- MIBs, page 42-13
- Application Services and Third-Party Tools, page 42-14

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 42-2 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 42-2 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.
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.

Table 42-2 Feature History for SNMP (continued)

Feature Name	Platform Releases	Feature Information
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**.
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.
SNMP OIDs	9.1(1)	Five new SNMP Physical Vendor Type OIDs have been added to support the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X.
NAT MIB	9.1(2)	Added the cnatAddrBindNumberOfEntries and cnatAddrBindSessionCount OIDs to support the xlate_count and max_xlate_count entries, which are the equivalent to allowing polling using the **show xlate count** command.
SNMP hosts, host groups, and user lists	9.1(5)	You can now add up to 4000 hosts. The number of supported active polling destinations is 128. You can specify a network object to indicate the individual hosts that you want to add as a host group. You can associate more than one user with one host. We modified the following screen: Configuration > Device Management > Management Access > SNMP.



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 43-1
- Licensing Requirements for NSEL, page 43-4
- Prerequisites for NSEL, page 43-4
- Guidelines and Limitations, page 43-4
- Configuring NSEL, page 43-5
- Monitoring NSEL, page 43-7
- Where to Go Next, page 43-7
- Additional References, page 43-7
- Feature History for NSEL, page 43-8

Information About NSEL

This section includes the following topics:

- Using NSEL and Syslog Messages, page 43-2
- Using NSEL in Clustering, page 43-3

The ASA and ASASM support NetFlow Version 9 services. For more information about NetFlow services, see the "RFCs" section on page 43-8.

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, and flow-denied (excluding those flows that are denied by EtherType ACLs). In addition, the ASA and ASASM implementation of NSEL generates periodic NSEL events, 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, they may also be triggered by state changes in the flow.

Note The flow-update event feature is not available in Version 9.0(1). It is available in Versions 8.4(5) and 9.1(2).

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.
- Triggers flow-update 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, flow-update, 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 ACL 1 to collector 1.
  - Log all flow-create events to collector 1.
  - Log all flow-teardown events to collector 2.
  - Log all flow-update events to collector 1.
- Delays the export of flow-create events.

Using NSEL and Syslog Messages

Table 43-1 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 43-1 redundant. In the interest of performance, we recommend that you disable redundant syslog messages, because the same information is exported through NetFlow.

Table 43-1 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). The fields in the header of the export packet include the system up time and UNIX time (synchronized across the cluster). These fields are all local to an individual ASA. The NSEL collector uses the combination of the source IP address and 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.
- If you have previously configured flow-export actions using the **flow-export event-type all** command, and you upgrade to a later version, NSEL automatically begins issuing flow-update records when necessary.
- 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.

Configuring NSEL

This section describes how to configure NSEL and includes the following topics:

- Using NetFlow, page 43-5
- Matching NetFlow Events to Configured Collectors, page 43-6

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 Enter the flow update interval, which specifies the time interval between flow-update events in minutes. Valid values are from 1 - 60 minutes. The default value is 1 minute.

Step 3 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 4 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 5 To configure more collectors, repeat Step 4 for each additional collector.

Step 6 To change collector configuration details, select a collector and click **Edit**. To remove a configured collector, select it and click **Delete**.

Step 7 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 8 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 43-6.

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 firewall configuration guide.

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 not available in Version 9.0(1). It is available in Versions 8.4(5) and 9.1(2).

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 4 of the "Using NetFlow" section on page 43-5.

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 43-7.

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 43-2.

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 Type **show logging flow-export-syslogs**, then press **Send**.	Lists all syslog messages that are captured by NSEL events.
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 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 41, "Configuring Logging."

Additional References

For additional information related to implementing NSEL, see the following sections:

- Related Documents, page 43-8
- RFCs, page 43-8

Related Documents

Related Topic	Document Title
Using NSEL and Syslog Messages, page 43-2	*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 43-2 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 43-2 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 ACLs. 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 43-2 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. **Note** The flow-update event feature is not available in Version 9.0(1).
NSEL	9.1(2)	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.



Configuring Anonymous Reporting and Smart Call Home

The Smart Call Home feature provides personalized, e-mail-based and web-based notification to you about critical events involving your individual systems, often before you 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.

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 44-1
- Licensing Requirements for Anonymous Reporting and Smart Call Home, page 44-3
- Prerequisites for Smart Call Home and Anonymous Reporting, page 44-4
- Guidelines and Limitations, page 44-4
- Configuring Anonymous Reporting and Smart Call Home, page 44-5
- Monitoring Anonymous Reporting and Smart Call Home, page 44-9
- Feature History for Anonymous Reporting and Smart Call Home, page 44-10

Information About Anonymous Reporting and Smart Call Home

This section includes the following topics:

- Information About Anonymous Reporting, page 44-1
- Information About Smart Call Home, page 44-3

Information About Anonymous Reporting

You 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

DNS Requirement

A DNS server must be configured correctly 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 that was 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 each Smart Call Home message that is 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 requests you to enable the Anonymous Reporting and Smart Call Home features according to the following guidelines:

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 44-5 or the "Configuring Smart Call Home" section on page 44-5.

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 you regarding your 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 prerequisite:

- DNS must be configured. See the "DNS Requirement" section on page 44-2 and the "Configuring the DNS Server" section on page 16-8.

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, you can remove the trust point, but disabling Anonymous Reporting does not cause the trust point to be removed.
- If you are using a multiple context mode configuration, the **dns**, **interface**, and **trustpoint** commands are in the admin context, and the **call-home** commands are in the system context.

Additional Guidelines for Smart Call Home

- In multiple context mode, the **subscribe-to-alert-group snapshot 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.
- A Smart Call Home message is sent to Cisco to report important cluster events if you have enabled clustering and configured Smart Call Home to subscribe to the diagnostic alert group with a critical severity level. A Smart Call Home clustering message is sent for only the following events:
  - When a unit joins the cluster
  - When a unit leaves the cluster
  - When a cluster unit becomes the cluster master
  - When a secondary unit fails in the cluster

  Each message that is sent includes the following information:
  - The active cluster member count

– The output of the **show cluster info** command and the **show cluster history** command on the cluster master

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 44-5
- Configuring Smart Call Home, page 44-5

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 44-9.)

Configuring Smart Call Home

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.

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 **Advanced System Setup**. This area 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, you 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, enter the 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.

Step 14 To determine if the configured alerts are operating correctly, click **Test**.

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 44-1 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 44-1 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	9.0(1)	You 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.
Smart Call Home	9.1(2)	The **show local-host** command was changed to the **show local-host | include interface** command for telemetry alert group reporting.
Smart Call Home	9.1(3)	A Smart Call Home message is sent to Cisco to report important cluster events if you have enabled clustering and configured Smart Call Home to subscribe to the Diagnostic alert group with a Critical severity level. A Smart Call Home clustering message is sent for only the following three events: • When a unit joins the cluster • When a unit leaves the cluster • When a cluster unit becomes the cluster master Each message that is sent includes the following information: • The active cluster member count • The output of the **show cluster info** command and the **show cluster history** command on the cluster master



PART 9

System Administration



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 45-1
- Configuring CLI Parameters, page 45-5
- Configuring ICMP Access, page 45-8
- Configuring Management Access Over a VPN Tunnel, page 45-10
- Configuring AAA for System Administrators, page 45-12
- Monitoring Device Access, page 45-32
- Feature History for Management Access, page 45-33

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 45-1
- Guidelines and Limitations, page 45-2
- Configuring Management Access, page 45-3
- Using a Telnet Client, page 45-4
- Using an SSH Client, page 45-5

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 45-10.
- 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.

- (9.1(2) and later) The default Telnet login password was removed; you must manually set the password before using Telnet. See the "Configuring the Hostname, Domain Name, and Passwords" section on page 16-1.
- 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 45-1.
- The AES-CTR encryption for SSH supports only AES-128 on single-core platforms, which include the ASA 5505, 5510, 5520, 5540, and 5550.

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—Specify 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.

c. DH Key Exchange—Click the applicable radio button to choose Diffie-Hellman (DH) Key Exchange Group 1 or Group 14. Both the DH Group 1 and Group 14 key-exchange methods for key exchange are supported on the ASA. If no DH group key-exchange method is specified, the DH group 1 key-exchange method is used. For more information about using DH key-exchange methods, see RFC 4253.

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.

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 33-3.

j. Click **OK,** then click **Apply**.

Using a Telnet Client

To gain access to the ASA CLI using Telnet, enter the login password. (9.1(2) and later) The default Telnet login password was removed; you must manually set the password before using Telnet. See the "Configuring the Hostname, Domain Name, and Passwords" section on page 16-1.

If you configure Telnet authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 45-17), 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:

```
ciscoasa(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.

You can alternatively configure a public key instead of using a password. See the "Adding a User Account to the Local Database" section on page 33-3.

Configuring CLI Parameters

This section includes the following topics:

- Licensing Requirements for CLI Parameters, page 45-5
- Guidelines and Limitations, page 45-5
- Configuring a Login Banner, page 45-6
- Customizing a CLI Prompt, page 45-7
- Changing the Console Timeout, page 45-8

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 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 6, "Configuring Access Rules," in the firewall configuration guide.

This section includes the following topics:

- Information About ICMP Access, page 45-9
- Licensing Requirements for ICMP Access, page 45-9
- Guidelines and Limitations, page 45-9
- Default Settings, page 45-9
- Configuring ICMP Access, page 45-10

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 45-11
- Guidelines and Limitations, page 45-2
- Configuring a Management Interface, page 45-11

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.

- Information About AAA for System Administrators, page 45-12
- Licensing Requirements for AAA for System Administrators, page 45-15
- Prerequisites, page 45-15
- Guidelines and Limitations, page 45-16
- Default Settings, page 45-17
- Configuring Authentication for CLI, ASDM, and enable command Access, page 45-17
- Limiting User CLI and ASDM Access with Management Authorization, page 45-18
- Configuring a Password Policy for Local Database Users, page 45-20
- Configuring Command Authorization, page 45-23
- Configuring Management Access Accounting, page 45-28
- Viewing the Currently Logged-In User, page 45-29
- Setting a Management Session Quota, page 45-30
- Recovering from a Lockout, page 45-31

Information About AAA for System Administrators

This section describes AAA for system administrators and includes the following topics:

- Information About Management Authentication, page 45-12
- Information About Command Authorization, page 45-13

Information About Management Authentication

This section describes authentication for management access and includes the following topics:

- Comparing CLI Access with and without Authentication, page 45-12
- Comparing ASDM Access with and without Authentication, page 45-13
- Authenticating Sessions from the Switch to the ASA Services Module, page 45-13

Comparing CLI Access with and without Authentication

How you log into the ASA depends on whether or not you enable authentication:

- No Authentication—If you do not enable any authentication for Telnet, you do not enter a username; you enter the login password. (SSH is not available without authentication). You access user EXEC mode.
- Authentication—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:

- No 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.
- 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 45-13
- About Preserving User Credentials, page 45-14
- Security Contexts and Command Authorization, page 45-14

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 45-23). (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

Prerequisites for the AAA Server or Local Database

You must configure users in a AAA server or the local database. For a AAA server, you then need to configure the ASA to communicate with it. See the following chapters:

- AAA server—See theapplicable AAA server-type chapter.
- Local Database—See the"Adding a User Account to the Local Database" section on page 33-3.

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 45-1.

Prerequisites for Local Command Authorization

- Configure **enable** authentication. (See the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 45-17.)

  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"Configuring LDAP Attribute Maps" section on page 36-5.

Prerequisites for TACACS+ Command Authorization

- Configure CLI and **enable** authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 45-17).

Prerequisites for Management Accounting

- Configure CLI and **enable** authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 45-17).

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 45-25.

Configuring Authentication for CLI, ASDM, and enable command Access

You can require authentication for CLI, ASDM, and enable command access.

Prerequisites

- Configure Telnet, SSH, or HTTP access according to the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 45-1.
- For SSH access, you must configure SSH authentication; there is no default username.

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 45-13.
- **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 45-13.

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

The ASA enables you to distinguish between administrative and remote-access users when they authenticate using RADIUS, LDAP, TACACS+, or the local user database. User role differentiation can prevent remote access VPN and network access users from establishing an administrative connection to the ASA.

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 Choose one of the following options:

- 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.

  When the **LOCAL** option is configured, the local user database is the source for the username entered and the Service-Type and Privilege-Level attributes assigned.

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 45-23 for more information.

When the **authentication-server** option is configured, the same server is used for both authentication and authorization.

- To enable management authorization, choose **Configuration > Device Management > Users/AAA > AAA Access > Authorization**, and check the **Allow privileged users to enter into EXEC mode on login** check box.

  The **auto-enable** option allows users with sufficient privileges from the login authentication server to be placed directly in privileged EXEC mode. Otherwise, users are placed in user EXEC mode. These privileges are determined by the Service-Type and Privilege-Level attributes that are required to enter each EXEC mode. To enter privileged EXEC mode, users must have a Service-Type attribute of Administrative and a Privilege Level attribute of greater than 1 assigned to them.

  This option is not supported in the system context. However, if you configure Telnet or serial authentication in the admin context, then authentication also applies to sessions from the switch to the ASASM.

  There is no effect if you enter the **aaa authorization exec** command alone.

  The **auto-enable** option is not included when you use serial authentication in management authorization.

  The **aaa authentication http** command is not affected by the **auto-enable** option.

  Before you configure the **auto-enable** option, we recommend that you configure both protocol login and enable authentication, and that all authentication requests go to the same AAA server group, as shown in the following example:

```
ciscoasa (config)# aaa authentication ssh console RADIUS
ciscoasa (config)# aaa authentication enable console RADIUS
ciscoasa (config)# aaa authorization exec authentication-server auto-enable
```

  We do *not* recommend that you use other types of configurations.

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

When users are authenticated through LDAP, the native LDAP attributes and their values can be mapped to Cisco ASA attributes to provide specific authorization features. Configure Cisco VSA CVPN3000-Privilege-Level with a value between 0 and 15. and then map the LDAP attributes to Cisco VAS CVPN3000-Privilege-Level. For more information, see the "Configuring LDAP Attribute Maps" section on page 36-5.

The RADIUS IETF **service-type** attribute, when sent in an access-accept message as the result of a RADIUS authentication and authorization request, is used to designate which type of service is granted to the authenticated user:

- 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. The Framed (2) and Login (1) service types are treated the same way.

- 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. All other service types (Voice, FAX, and so on) are treated the same way.

The RADIUS Cisco VSA **privilege-level** attribute (Vendor ID 3076, sub-ID 220), when sent in an access-accept message, is used to designate the level of privilege for the user.

When an authenticated user tries administrative access to the ASA through ASDM, SSH, or Telnet, but does not have the appropriate privilege level to do so, the ASA generates syslog message 113021. This message informs the user that the attempted login failed because of inappropriate administrative privileges.

TACACS+ users

Authorization is requested with "service=shell," and the server responds with PASS or FAIL.

- PASS, privilege level 1— Allows full access to any services specified by the Authentication tab options.
- PASS, privilege level 2 and higher—Allows access to the CLI when you configure the Telnet or SSH authentication options, but denies ASDM configuration access if you configure the HTTP option. ASDM monitoring access is allowed. If you configure **enable** authentication with the Enable option, the user cannot access privileged EXEC mode using the **enable** command. You are not allowed to access privileged EXEC mode using the **enable** command if your enable privilege level is set to 14 or less.
- FAIL—Denies management access. You 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 for a given username. 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 33-3.

Configuring a Password Policy for Local Database Users

When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters.

The password policy only applies to administrative users using the local database, and not to other types of traffic that can use the local database, such as VPN or AAA for network access, and not to users authenticated by a AAA server.

This section includes the following topics:

- Configuring the Password Policy, page 45-21
- Changing Your Password, page 45-22

Configuring the Password Policy

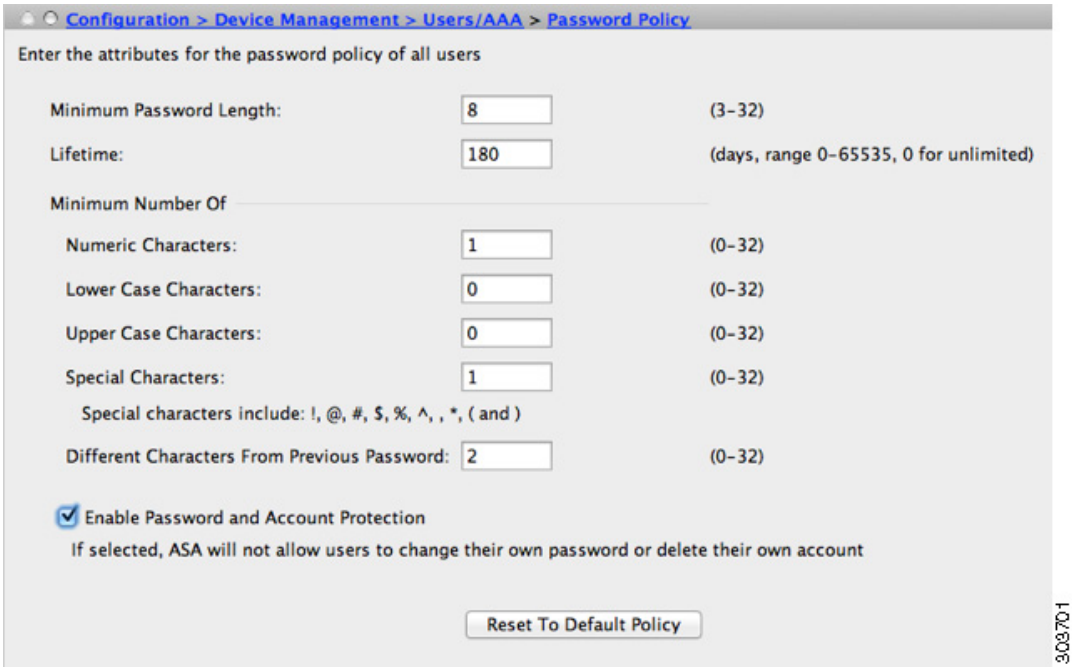
After you configure the password policy, when you change a password (either your own or another user's), the password policy applies to the new password. Any existing passwords are grandfathered in. The new policy applies to changing the password with the User Accounts pane as well as the Change My Password pane.

Prerequisites

- Configure both CLI/ASDM and enable authentication according to the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 45-17. Be sure to specify the local database.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Users/AAA > Password Policy**.



Step 2 Configure any mix of the following options:

- Minimum Password Length—Enter the minimum length for passwords. Valid values range from 3 to 64 characters. The recommended minimum password length is 8 characters.
- Lifetime—Enter the interval in days after which passwords expire for remote users (SSH, Telnet, HTTP); users at the console port are never locked out due to password expiration. Valid values are between 0 and 65536 days. The default value is 0 days, a value indicating that passwords will never expire.

  7 days before the password expires, a warning message appears. After the password expires, system access is denied to remote users. To gain access after expiration, do one of the following:

  - Have another administrator change your password.
  - Log in to the physical console port to change your password.

- Minimum Number Of—Specify the minimum of characters from the following types:
  - Numeric Characters—Enter the minimum number of numeric characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0.
  - Lower Case Characters—Enter the minimum number of lower case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.
  - Upper Case Characters—Enter the minimum number of upper case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.
  - Special Characters—Enter the minimum number of special characters that passwords must have. Valid values range from 0 to 64 characters. Special characters include the following: !, @, #, $, %, ^, &, *, '(‘ and ‘)’.The default value is 0.
  - Different Characters from Previous Password—Enter the minimum number of characters that you must change between new and old passwords. Valid values are between 0 and 64 characters. The default value is 0. Character matching is position independent, meaning that new password characters are considered changed only if they do not appear anywhere in the current password.

Step 3 (Optional) Check the **Authentication Enable** check box to require users to change their password on the Change My Password pane instead of the User Accounts pane. The default setting is disabled: a user can use either method to change their password.

If you enable this feature, if you try to change your password on the User Accounts pane, the following error message is generated:

```
ERROR: Changing your own password is prohibited
```

Step 4 To reset the password policy to the default, click **Reset to Default**.

Step 5 Click **Apply** to apply the configuration settings.

Changing Your Password

If you configure a password lifetime in the password policy, you need to change your password to a new one when the old password expires. This password change method is required if you enable password policy authentication. If password policy authentication is not enabled, then you can use this method, or you can change your user account directly with the User Accounts pane.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Users/AAA > Change Password**.



Step 2 Enter your old password.

Step 3 Enter your new password.

Step 4 Confirm your new password.

Step 5 Click **Make Change**.

Step 6 Click the **Save** icon to save your changes to the running configuration.

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 45-13.

This section includes the following topics:

- Configuring Local Command Authorization, page 45-23
- Viewing Local Command Privilege Levels, page 45-25
- Configuring Commands on the TACACS+ Server, page 45-25
- Configuring TACACS+ Command Authorization, page 45-27

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 following sections for more information:

- "Adding a User Account to the Local Database" section on page 33-3
- "Supported Authentication Methods" section on page 34-2
- "Configuring LDAP Attribute Maps" section on page 36-5

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 privilege level of a command, double-click it or click **Edit** to display the Select Command Privilege dialog box. You can set the privilege level between 0 and 15. You can only configure the privilege level of the *main* command. For example, you can configure the privilege level of all **aaa** commands, but not the privilege level of the **aaa authentication** command and the **aaa authorization** command separately.

  To change the privilege 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, mapped LDAP, and TACACS+ users. See the "Limiting User CLI and ASDM Access with Management Authorization" section on page 45-18 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.

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.

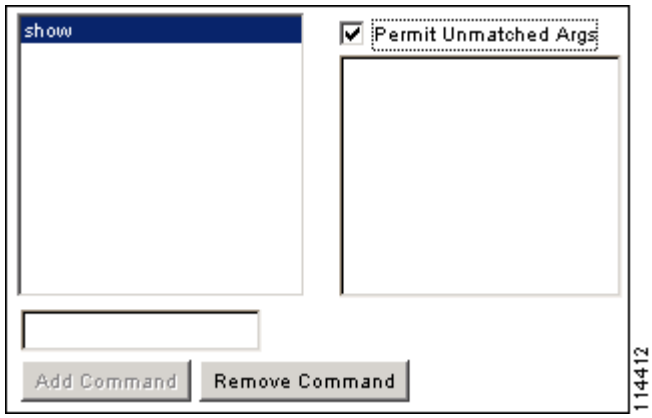
- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

  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.

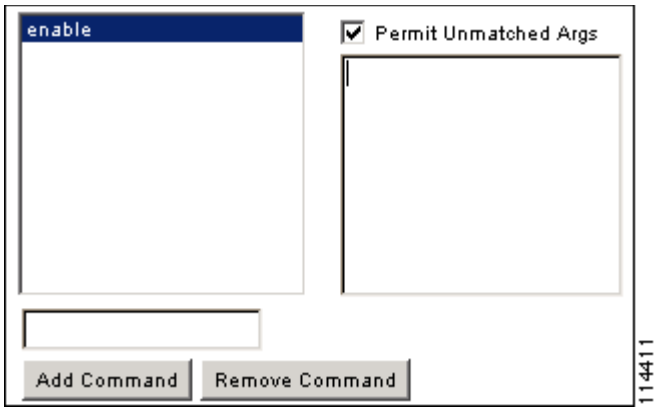
  For example, you can configure just the **show** command, then all the **show** commands are allowed. We recommend using this method so that you do not have to anticipate every variant of a command, including abbreviations and a question mark, which shows CLI usage (see Figure 45-1).

Figure 45-1 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 45-2).

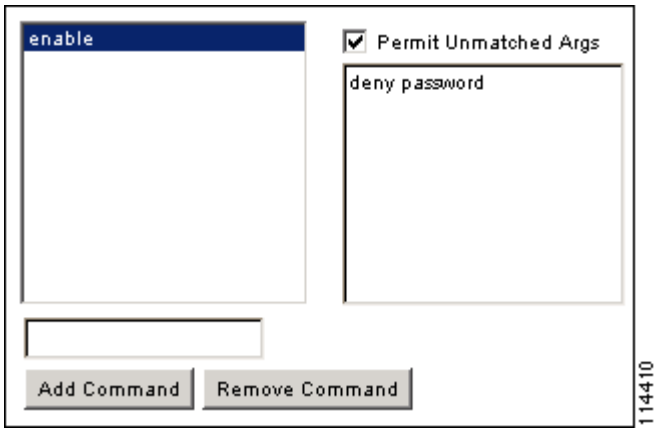
Figure 45-2 Permitting Single Word Commands



- 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 45-3).

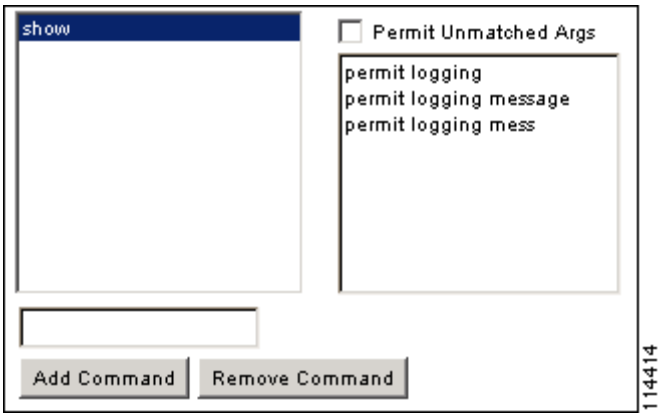
Figure 45-3 Disallowing Arguments



- 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 45-4).

Figure 45-4 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 45-31.

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 45-23.

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 33-3) and command privilege levels (see the "Configuring Local Command Authorization" section on page 45-23).

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, in the Tools > Command Line Interface tool:

```
ciscoasa# show curpriv
```

Examples

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

```
ciscoasa# show curpriv
Username: admin
Current privilege level: 15
Current Mode/s: P_PRIV
```

Table 45-1 describes the **show curpriv** command output.

Table 45-1 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 Modes	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

Setting a Management Session Quota

You can establish a maximum number of simultaneous management sessions. If the maximum is reached, no additional sessions are allowed and a syslog message is generated. To prevent a system lockout, the management session quota mechanism cannot block a console session.

To set a management session quota, perform the following steps:

Step 1 Choose **Configuration > Device Management > Management Access > Management Session Quota**.

Step 2 Enter the maximum number of simultaneous ASDM, SSH, and Telnet sessions that are allowed on the ASA. Valid values range from 0 to 10000.



Note If the management quota session number is exceeded, an error message appears, and ASDM closes.

Step 3 Click **Apply** to save the configuration changes.

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 45-2 lists the common lockout conditions and how you might recover from them.

Table 45-2 CLI Authentication and Command Authorization Lockout Scenarios

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
Local CLI authentication	No users have been configured 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	The server is 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 that you do not get locked out when the server is down.
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, outgoing HMACs, SSH session IDs, remaining rekey data, remaining rekey time, data-based rekeys, time-based rekeys, and the last rekey time. 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 45-3 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 45-3 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.

Table 45-3 Feature History for Management Access (continued)

Feature Name	Platform Releases	Feature Information
Support for administrator password policy when using the local database	8.4(4.1), 9.1(2)	When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters. We introduced the following screen: Configuration > Device Management > Users/AAA > Password Policy.
Support for SSH public key authentication	8.4(4.1), 9.1(2)	You can enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF. *PKF key format support is only in 9.1(2) and later.*
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	8.4(4.1), 9.1(2)	Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported. We modified the following screen: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH.
Support for a maximum number of management sessions	8.4(4.1), 9.1(2)	You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions. We introduced the following screen: Configuration > Device Management > Management Access > Management Session Quota.
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.
AES-CTR encryption for SSH	9.1(2)	The SSH server implementation in the ASA now supports AES-CTR mode encryption.
Improved SSH rekey interval	9.1(2)	An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic. .

Table 45-3 Feature History for Management Access (continued)

Feature Name	Platform Releases	Feature Information
Improved one-time password authentication	9.1(5)	Administrators who have sufficient authorization privileges may enter privileged EXEC mode by entering their authentication credentials once. The **auto-enable** option was added to the **aaa authorization exec** command. We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization.



Managing Software and Configurations

This chapter describes how to manage the ASA software and configurations and includes the following sections:

- Upgrading the Software, page 46-1
- Managing Files, page 46-12
- Configuring the Images and Startup Configuration to Use, page 46-19
- Backing Up and Restoring Configurations or Other Files, page 46-21
- Saving the Running Configuration to a TFTP Server, page 46-28
- Scheduling a System Restart, page 46-28
- Downgrading Your Software, page 46-29
- Configuring Auto Update, page 46-31

Upgrading the Software

This section describes how to upgrade to the latest version and includes the following topics:

- Upgrade Path and Migrations, page 46-1
- Viewing Your Current Version, page 46-3
- Downloading the Software from Cisco.com, page 46-3
- Upgrading a Standalone Unit, page 46-3
- Upgrading a Failover Pair or ASA Cluster, page 46-6

Note For CLI procedures, see the ASA documentation.

Upgrade Path and Migrations

- If you are upgrading from a pre-8.3 release:
  - See the *Cisco ASA 5500 Migration Guide to Version 8.3 and Later* for important information about migrating your configuration.
  - You cannot upgrade directly to 9.0 or later. You must first upgrade to Version 8.3 or 8.4 for a successful migration.

- If you are upgrading from a pre-9.0 release, because of ACL migration, you cannot later perform a downgrade; be sure to back up your configuration file in case you want to downgrade. See the ACL migration section in the 9.0 release notes for more information.
- If you are upgrading from one of the following versions, you can successfully upgrade to 9.1(2.8) and 9.1(3) or later:
  - 8.4(5) or later
  - 9.0(2) or later
  - 9.1(2)

  However, if you are running any earlier versions, you cannot upgrade directly to 9.1(2.8) or 9.1(3) or later without *first* upgrading to one of the above versions. For example:

ASA Version	First Upgrade to:	Then Upgrade to:
8.2(1)	8.4(6)	9.1(2.8) or 9.1(3) or later
8.4(4)	8.4(6)	9.1(2.8) or 9.1(3) or later
9.0(1)	9.0(3)	9.1(2.8) or 9.1(3) or later
9.1(1)	9.1(2)	9.1(2.8) or 9.1(3) or later

- 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 46-1 shows the supported scenarios for performing zero-downtime upgrades.

Table 46-1 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 8.4(1) to 8.4(6) without first installing the maintenance releases in between.

Table 46-1 Zero-Downtime Upgrade Support (continued)

Type of Upgrade	Support
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 8.2 to 8.3. Upgrading from 8.2 directly to 8.4 is not supported for zero-downtime upgrades; you must first upgrade to 8.3. For models that are not supported on a minor release, you can skip the minor release; for example, for the ASA 5585-X, you can upgrade from 8.2 to 8.4 (the model is not supported on 8.3). **Note** Zero-downtime upgrades are possible, even when feature configuration is migrated, for example, from 8.2 to 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 8.6 to 9.0, assuming that 8.6 is the last minor version in the 8.x release series for your model. Upgrading from 8.6 directly to 9.1 is not supported for zero-downtime upgrades; you must first upgrade to 9.0. For models that are not supported on a minor release, you can skip the minor release; for example, for the ASA 5585-X, you can upgrade from 8.4 to 9.0 (the model is not supported on 8.5 or 8.6). **Note** Zero-downtime upgrades are possible, even when feature configuration is migrated, for example, from 8.4 to 9.0.

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/go/asa-software

Upgrading a Standalone Unit

This section describes how to install the ASDM and operating system (OS) images.

- Upgrading from Your Local Computer, page 46-4
- Upgrading Using the Cisco.com Wizard, page 46-5

Upgrading from Your Local Computer

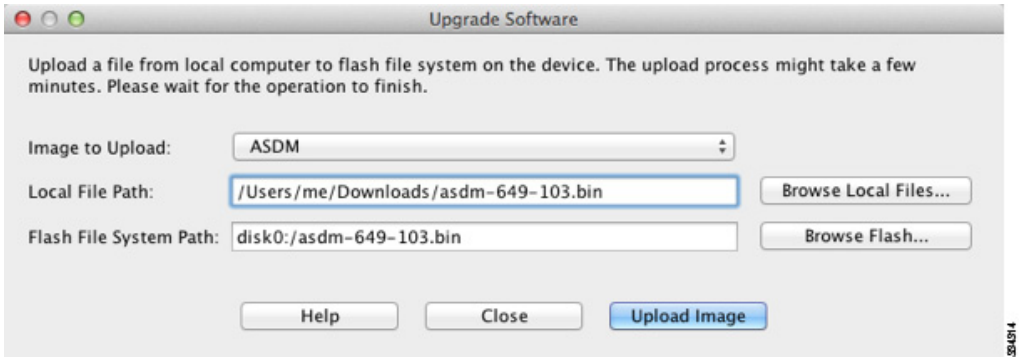
The Upgrade Software from Local Computer tool lets you upload an image file from your computer to the flash file system to upgrade the ASA.

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 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.

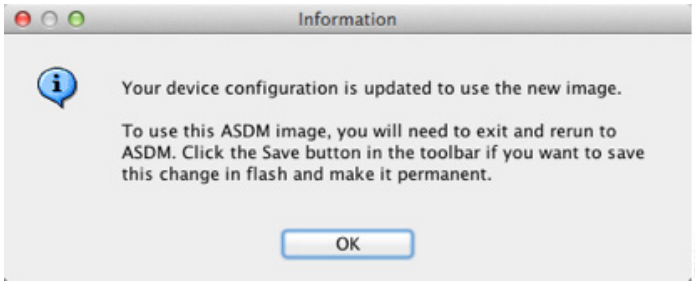
Step 5 In the Flash File System Path field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 You are prompted to set this image as the ASDM image. Click **Yes**.



Step 8 You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the Upgrade tool. **Note:** You will save the configuration and reload ASDM *after* you upgrade the ASA software.



Step 9 Repeat Step 2 through Step 8, choosing **ASA** from the Image to Upload drop-down list. You can also use this procedure to upload other file types.

Step 10 Choose **Tools > System Reload** to reload the ASA.

A new window appears that asks you to verify the details of the reload.

a. Click the **Save the running configuration at the time of reload** radio button (the default).

b. Choose a time to reload (for example, **Now**, the default).

c. 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

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 9.1(4), the download might be 9.1(4.2). 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 is 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).

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 a Failover Pair or ASA Cluster

- Upgrading an Active/Standby Failover Pair, page 46-6
- Upgrading an Active/Active Failover Pair, page 46-8
- Upgrading an ASA Cluster, page 46-10

Upgrading an Active/Standby Failover Pair

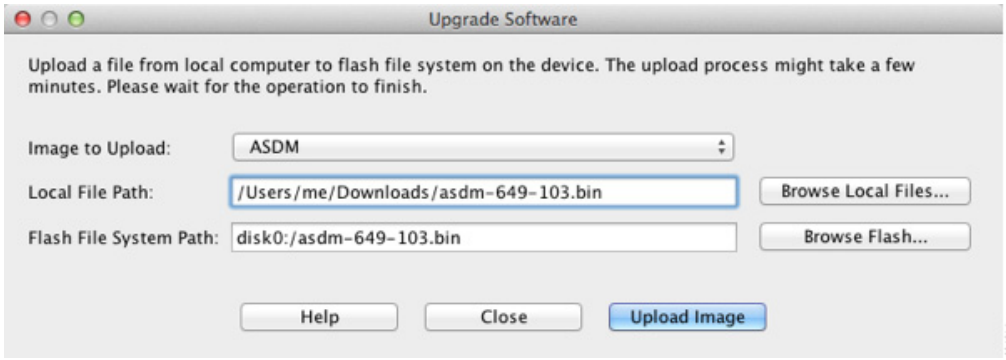
To upgrade the Active/Standby failover pair, perform the following steps.

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.

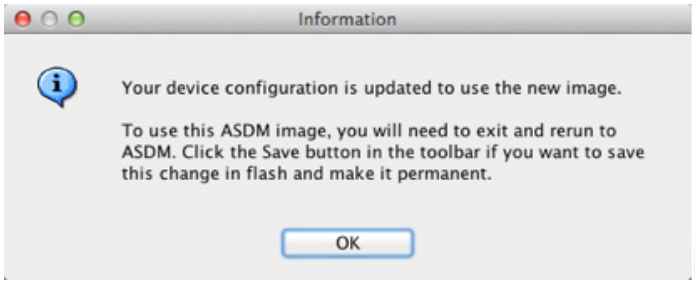
Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 You are prompted to set this image as the ASDM image. Click **Yes**.



Step 8 You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the Upgrade tool. **Note:** You will save the configuration and reload ASDM *after* you upgrade the ASA software.

Step 9 Repeat Step 2 through Step 8, choosing **ASA** from the Image to Upload drop-down list.



Step 10 Click the **Save** icon on the toolbar to save your configuration changes.

Step 11 Connect ASDM to the *standby* unit, and upload the ASA and ASDM software according to Step 2 through Step 9, using the same file locations you used on the active unit.

Step 12 Choose **Tools > System Reload** to reload the standby ASA.

A new window appears that asks you to verify the details of the reload.

a. Click the **Save the running configuration at the time of reload** radio button (the default).

b. Choose a time to reload (for example, **Now**, the default).

c. 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 After the standby ASA reloads, restart ASDM and connect to the standby unit to make sure it is running.

Step 14 Connect ASDM to the *active* unit again.

Step 15 Force the active unit to fail over to the standby unit by choosing **Monitoring > Properties > Failover > Status**, and clicking **Make Standby**.

Step 16 Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload.

a. Click the **Save the running configuration at the time of reload** radio button (the default).

b. Choose a time to reload (for example, **Now**, the default).

c. 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

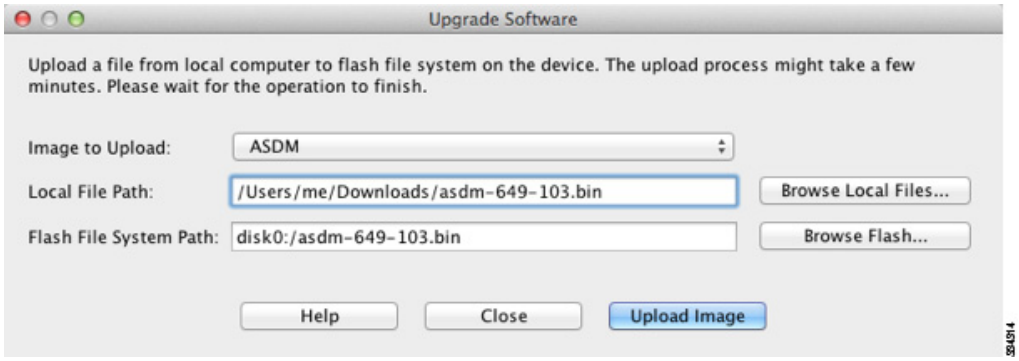
Perform these steps in the system execution space.

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.

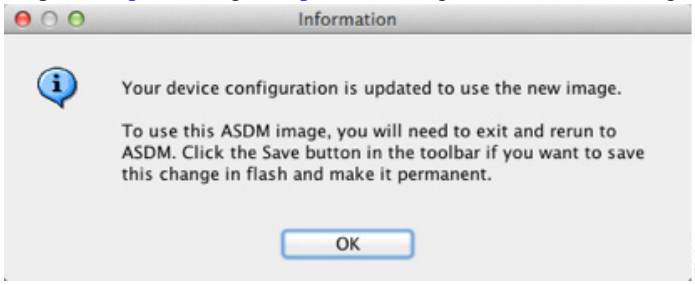
Step 6 Click **Upload Image**. The uploading process might take a few minutes.

Step 7 You are prompted to set this image as the ASDM image. Click **Yes**.



Step 8 You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the Upgrade tool. **Note:** You will save the configuration and reload ASDM *after* you upgrade the ASA software.

Step 9 Repeat Step 2 through Step 8, choosing **ASA** from the Image to Upload drop-down list.



Step 10 Click the **Save** icon on the toolbar to save your configuration changes.

Step 11 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 12 Connect ASDM to the *secondary* unit, and upload the ASA and ASDM software according to Step 2 through Step 9, using the same file locations you used on the active unit.

Step 13 Choose **Tools > System Reload** to reload the secondary ASA.

A new window appears that asks you to verify the details of the reload.

a. Click the **Save the running configuration at the time of reload** radio button (the default).

b. Choose a time to reload (for example, **Now**, the default).

c. 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 14 Connect ASDM to the *primary* unit, and check when the secondary unit reloads by choosing **Monitoring > Failover > System**.

Step 15 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 16 Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload.

a. Click the **Save the running configuration at the time of reload** radio button (the default).

b. Choose a time to reload (for example, **Now**, the default).

c. 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.

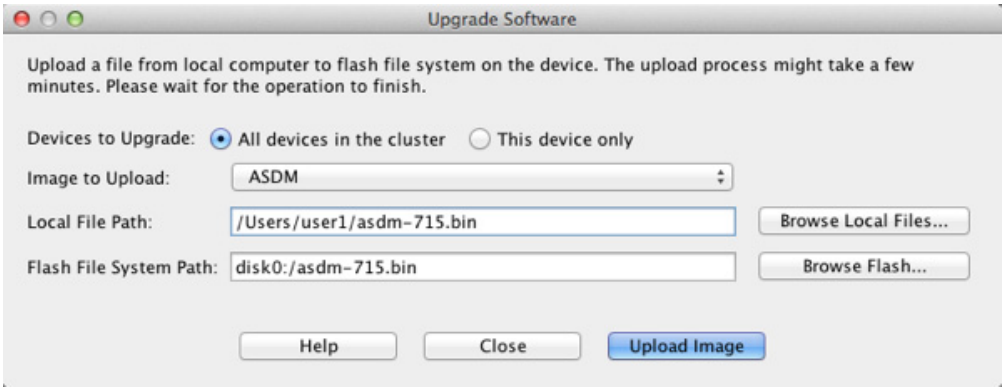
Step 2 (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.

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.

The Upgrade Software dialog box appears.



Step 5 From the Image to Upload drop-down list, choose **ASDM**.

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.

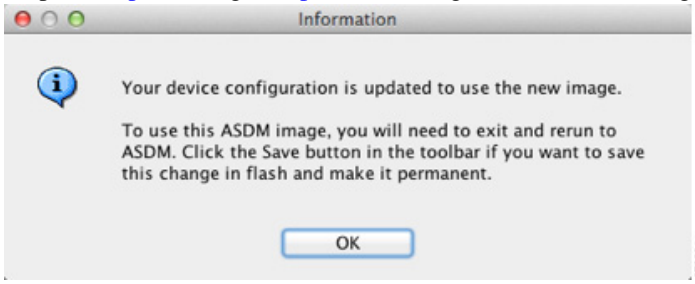
Step 8 Click **Upload Image**. The uploading process might take a few minutes.

Step 9 You are prompted to set this image as the ASDM image. Click **Yes**.



Step 10 You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the Upgrade tool. **Note:** You will save the configuration and reload ASDM *after* you upgrade the ASA software.

Step 11 Repeat Step 3 through Step 10, choosing **ASA** from the Image to Upload drop-down list.

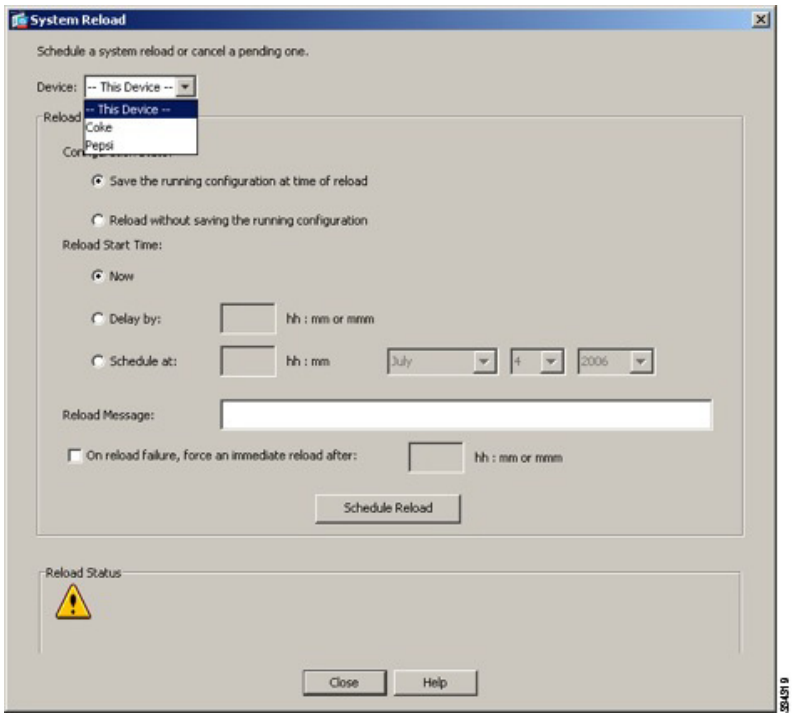


Step 12 Click the **Save** icon on the toolbar to save your configuration changes.

Step 13 Choose **Tools > System Reload**.

The System Reload dialog box appears.

Step 14 Reload each slave unit one at a time by choosing a slave unit name from the Device drop-down list, and then clicking **Schedule Reload** to reload the unit now.



To avoid connection loss and allow traffic to stabilize, wait for each unit to come back up (approximately 5 minutes) before reloading the next unit. To view when a unit rejoins the cluster, see the Monitoring > ASA Cluster > Cluster Summary pane.

Step 15 After all slave units have reloaded, disable clustering on the master unit by choosing **Configuration > Device Management > High Availability and Scalability > ASA Cluster**, uncheck the **Participate in ASA cluster** check box, and click **Apply**.

Wait for 5 minutes for a new master to be selected and traffic to stabilize. When the former master unit rejoins the cluster, it will be a slave.

Do not save the configuration; when the master unit reloads, you want clustering to be enabled on it.

Step 16 Choose **Tools > System Reload** and reload the master unit from the System Reload dialog box by choosing **--This Device--** from the Device drop-down list.

Step 17 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.

- Configuring File Access, page 46-13

- Accessing the File Management Tool, page 46-17
- Transferring Files, page 46-18

Configuring File Access

- Configuring the FTP Client Mode, page 46-13
- Configuring the ASA as a Secure Copy Server, page 46-13
- Customizing the ASA Secure Copy Client, page 46-14
- Configuring the ASA TFTP Client Path, page 46-15
- Adding Mount Points, page 46-15

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.

Detailed 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 (SCP) server on the ASA. Only clients that are allowed to access the ASA using SSH can establish a secure copy connection.

Restrictions

- 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.

Prerequisites

- Enable SSH on the ASA according to the "Configuring ASA Access for ASDM, Telnet, or SSH" section on page 45-1.
- The ASA license must have the strong encryption (3DES/AES) license to support SSH Version 2 connections.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server**, and check the **Enable secure copy server** check box.

Step 2 Click **Apply**.

Example

From a client on the external host, perform an SCP file transfer. For example, in Linux enter the following command:

```
scp -v -pw password source_filename username@asa_address:{disk0|disk1}:/dest_filename
```

The **-v** is for verbose, and if **-pw** is not specified, you will be prompted for a password.

Customizing the ASA Secure Copy Client

You can copy files to and from the ASA using the on-board SCP client (see the "Accessing the File Management Tool" section on page 46-17). This section lets you customize the SCP client operation.

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 **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP)**.
- For multiple mode in the System, choose **Configuration > Device Management > Device Administration > Secure Copy**

Step 2 The ASA stores the SSH host key for each SCP server to which it connects. You can manually add or delete servers and their keys from the ASA database if desired.

To add a key:

a. Click **Add** for a new server, or select the server from the Trusted SSH Hosts table, and click **Edit**.

b. For a new server, in the Host field, enter the server IP address.

c. Check the **Add public key for the trusted SSH host** check box.

d. Specify one of the following keys:

- Fingerprint—Enter the already hashed key; for example, a key that you copied from **show** command output.
- Key—Enter the public key or hashed value of the SSH host. The key string is the Base64 encoded RSA public key of the remote peer. You can obtain the public key value from an open SSH client; that is, from the .ssh/id_rsa.pub file. After you submit the Base64 encoded public key, that key is then hashed via SHA-256.

To delete a key:

a. Select the server from the Trusted SSH Hosts table, and click **Delete**.

Step 3 To be informed when a new host key is detected, check the **Inform me when a new host key is detected** check box.

By default, this option is enabled. When this option is enabled, you are prompted to accept or reject the host key if it is not already stored on the ASA. When this option is disabled, the ASA accepts the host key automatically if it was not stored before.

Step 4 Click **Apply**.

Configuring the ASA TFTP Client Path

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 copy files to or from a TFTP *server* (see the "Transferring Files" section on page 46-18. In this way, you can back up and propagate configuration files to multiple ASAs.

This section lets you pre-define the path to a TFTP server so you do not need to enter it in commands such as **copy** and **configure net**.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Management Access > File Access > TFTP Client**, and 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**.

Adding Mount Points

This section includes the following topics:

- Adding a CIFS Mount Point, page 46-15
- Adding an FTP Mount Point, page 46-16

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.

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 46-18 for more information.

Step 10 Click **Mount Points** to open the Manage Mount Points dialog box. See the "Adding Mount Points" section on page 46-15 for more information.

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, SMB, or SCP.

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 46-18
- Transferring Files Between Remote Server and Flash, page 46-18

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 46-20.
- 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 46-21
- Backing Up the Local CA Server, page 46-24
- Restoring Configurations, page 46-25
- Saving the Running Configuration to a TFTP Server, page 46-28

Backing Up Configurations

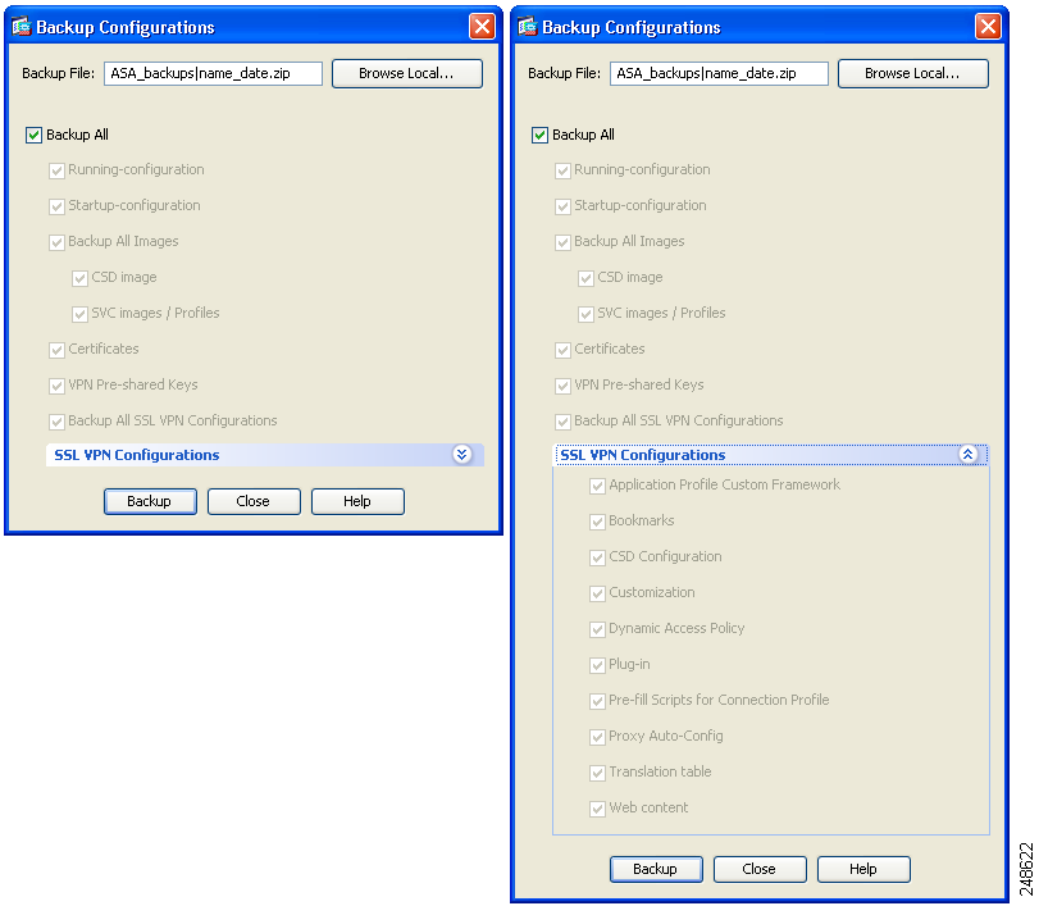
This procedure explains how to back up configurations and images to a .zip file and transfer it to your local computer.

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 16-5 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.**

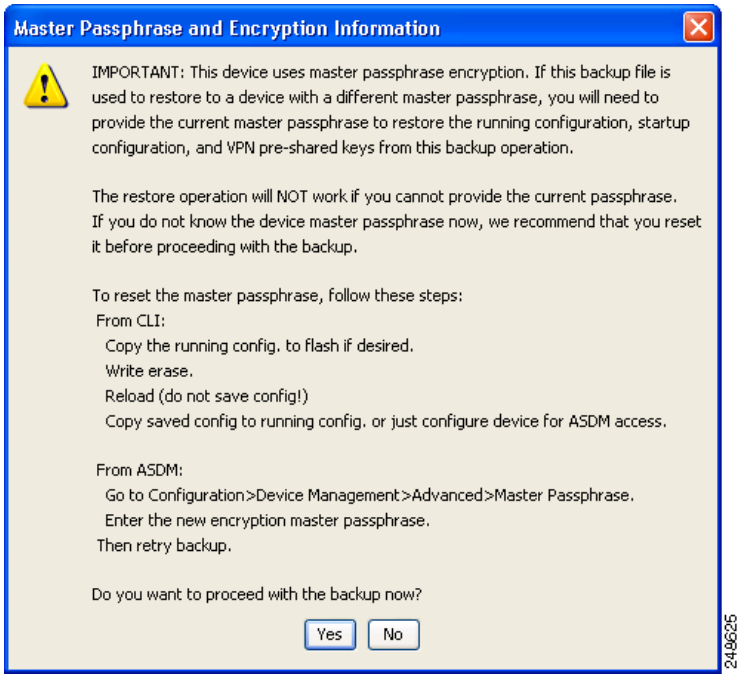
Step 6 In the Select dialog box, choose the directory in which you want to store the backup file.

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.

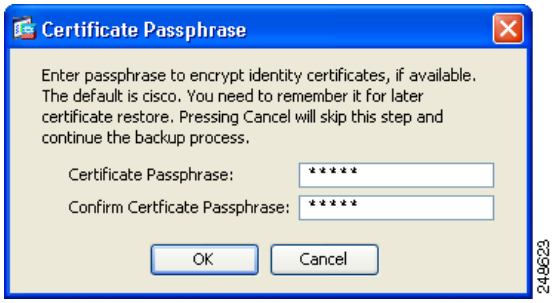
Step 9 Click **Backup**. The backup proceeds immediately unless you are backing up certificates or the ASA is using a master passphrase.

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 46-24.

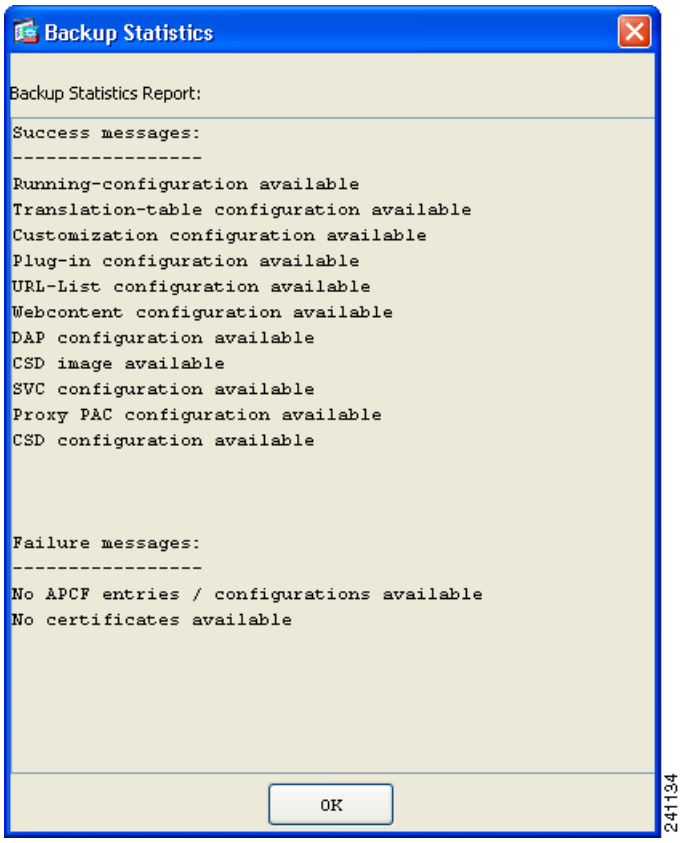


- To encrypt certificates, enter and confirm your certificate passphrase in the Certificate Passphrase dialog box and click **OK**. You will need to remember the password you enter in this dialog box when restoring the certificates.
- Clicking **Cancel** skips the step and does not back up certificates.

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:

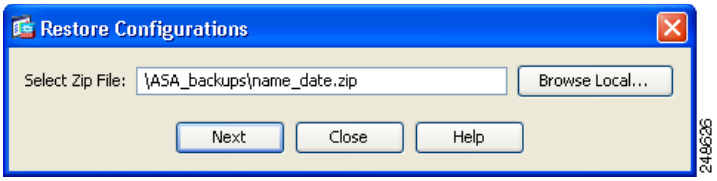
- The zip file that you restore must be created by choosing the Tools > Backup Configurations option.
- If you performed the backup with the master passphrase enabled, then you will need that master passphrase in order to restore the running configuration, start-up configuration, and VPN pre-shared keys 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 16-5 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.
- If you import PKCS12 data (with the **crypto ca trustpoint** command) and the trustpoint uses RSA keys, the imported key pair is assigned the same name as the trustpoint. Because of this limitation, if you specify a different name for the trustpoint and its key pair after you have restored an ASDM configuration, the startup configuration will be the same as the original configuration, but the running configuration will include a different key pair name. This means that if you use different names for the key pair and trustpoint, you cannot restore the original configuration. To work around this issue, make sure that you use the same name for the trustpoint and its key pair.

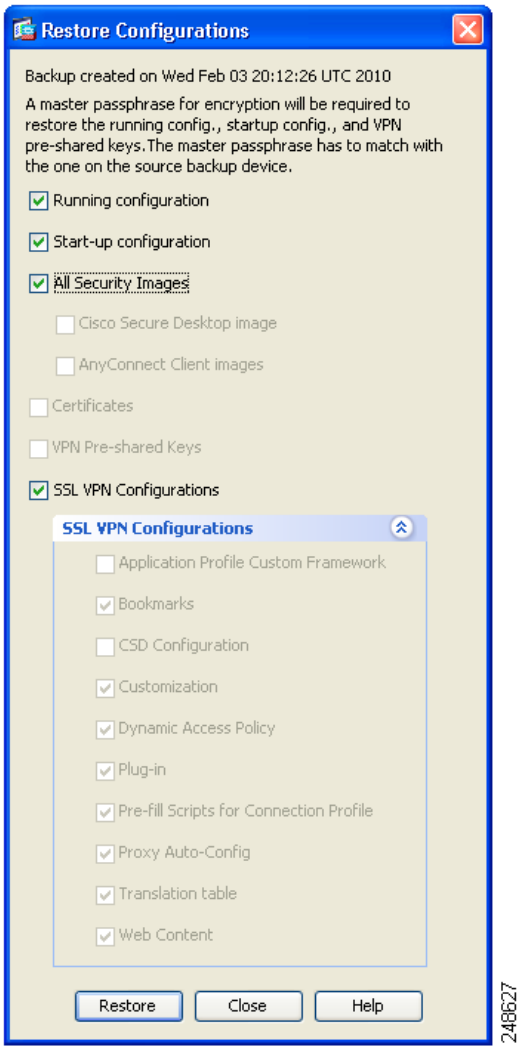
To restore selected elements of the ASA configuration, Cisco Secure Desktop image, or SSL VPN Client images and profiles, perform the following steps:

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 46-29
- Performing the Downgrade, page 46-30

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 46-21 for more information about configuration migration.

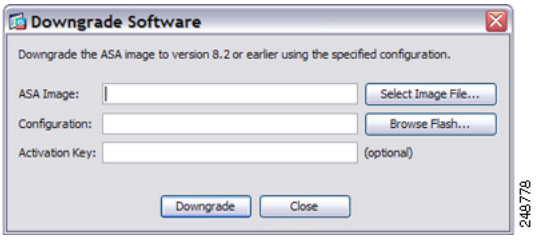
To downgrade from Version 8.3, perform the following steps:

Detailed Steps

Step 1 Choose **Tools > Downgrade Software**.

The Downgrade Software dialog box appears.

Figure 46-1 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 46-29 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 46-31
- Guidelines and Limitations, page 46-34
- Configuring Communication with an Auto Update Server, page 46-34

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.

- Auto Update Client or Server, page 46-31
- Auto Update Benefits, page 46-31
- Auto Update Server Support in Failover Configurations, page 46-32

Auto Update Client or Server

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 Benefits

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.

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
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.

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.

Feature History for Software and Configurations

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 Software and Configurations

Feature Name	Platform Releases	Feature Information
Secure Copy client	9.1(5)	The ASA now supports the Secure Copy (SCP) client to transfer files to and from a SCP server. We modified the following screens: Tools > File Management > File Transfer > Between Remote Server and Flash **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) 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 47-1

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 ACLs 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** to display the Ingress Selector screen.

Step 4 To capture packets on an interface, click the **Select Interface** radio button. To capture packets on the ASA CX dataplane, click the **Use backplane channel** radio button.

Step 5 In the Packet Match Criteria area, do one of the following:

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

Step 6 Tocontinue, see the "Ingress Traffic Selector" section on page 47-3.

Step 7 Click **Next** to display the Egress Traffic Selector screen. To continue, see the "Egress Traffic Selector" section on page 47-4.

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 8 Click **Next** to display the Buffers & Captures screen. To continue, see the "Buffers" section on page 47-4.

Step 9 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 10 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 11 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. To continue, see the "Summary" section on page 47-4.

Step 12 Click **Next** to display the Run Captures screen, and then click **Start** to begin capturing packets. Click **Stop** to end the capture. To continue, see the "Run Captures" section on page 47-4. If you are using clustering, go to Step 14.

Step 13 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 14 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 15 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. To continue, see the "Save Captures" section on page 47-5.

Step 16 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 17 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the ingress capture.

Step 18 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 19 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the egress capture.

Step 20 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 In the Point of Ingress area, choose the ingress interface name from the drop-down list.

Step 2 Enter the ingress source host and network. To capture packets on the ASA CX dataplane, click the **Use backplane channel** radio button.

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 To capture packets on an interface, click the **Select Interface** radio button. To capture packets on the ASA CX dataplane, click the **Use backplane channel** radio button.

Step 2 In the Point of Egress area, choose the egress interface name from the drop-down list.

Step 3 Enter the egress source host and network.

Step 4 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 10

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 48-1
- IPv6 Addresses, page 48-5
- Protocols and Applications, page 48-11
- TCP and UDP Ports, page 48-11
- Local Ports and Protocols, page 48-14
- ICMP Types, page 48-15

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 48-1
- Private Networks, page 48-2
- Subnet Masks, page 48-2

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 48-3
- Determining the Address to Use with the Subnet Mask, page 48-3

Determining the Subnet Mask

To determine the subnet mask based on how many hosts you want, see Table 48-1.

Table 48-1 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 48-3
- Class B-Size Network Address, page 48-4

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 48-2 shows the 8-host subnets (/29) of 192.168.0.x.

Table 48-2 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 48-2 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 48-3 shows the 16 subnets of the network 10.1.

Table 48-3 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 48-5
- IPv6 Address Types, page 48-6
- IPv6 Address Prefixes, page 48-10

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 13-15

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 48-4 shows several examples of address compression for different types of IPv6 address.

Table 48-4 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 48-6
- Multicast Address, page 48-8
- Anycast Address, page 48-9
- Required Addresses, page 48-10

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 48-7
- Site-Local Address, page 48-7
- Link-Local Address, page 48-7
- IPv4-Compatible IPv6 Addresses, page 48-7
- Unspecified Address, page 48-8
- Loopback Address, page 48-8
- Interface Identifiers, page 48-8

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 48-10, 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 48-8, 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 48-7).

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.

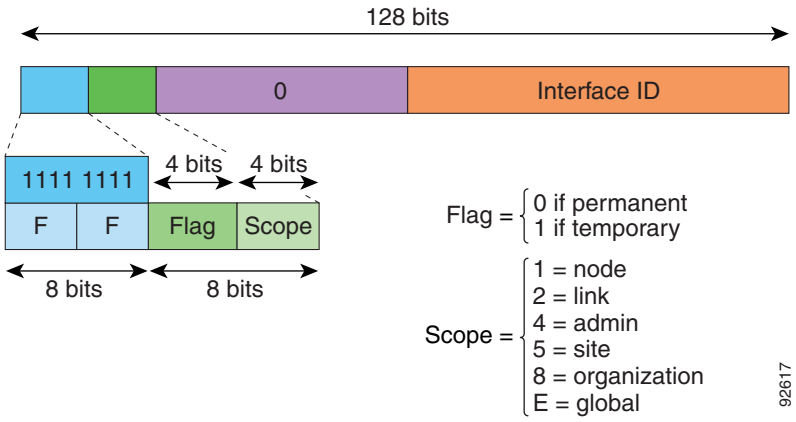
For example, and interface with the MAC address of 00E0.b601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

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 48-1 shows the format of the IPv6 multicast address.

Figure 48-1 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 48-5 shows the prefixes for each IPv6 address type.

Table 48-5 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 48-6 lists the protocol literal values and port numbers; either can be entered in ASA commands.

Table 48-6 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 48-7 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 48-7 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 48-7 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 48-7 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 48-8 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 48-8, 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 48-8 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 48-8 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 48-9 lists the ICMP type numbers and names that you can enter in ASA commands.

Table 48-9 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 48-9 ICMP Types (continued)

ICMP Number	ICMP Name
31	conversion-error
32	mobile-redirect



INDEX

Symbols

/bits subnet masks 48-3

Numerics

4GE SSM
  connector types 11-14
  fiber 11-14
  SFP 11-14
802.1Q tagging 12-10
802.1Q trunk 11-36

A

AAA
  about 32-1, 33-1, 34-1, 35-1, 37-1
  authentication
    CLI access 45-17
  authorization
    command 45-23
  server 41-4
    adding 34-15, 36-7, 36-8, 37-3
    types 32-1
  support summary 32-3
ABR
  definition of 27-2
access_rules 22-3
Access Group pane
  description 30-10
access ports 12-8
ACL Manager
  dialog box 21-1
activation key
  entering 5-37
  location 5-35
  obtaining 5-36
Active/Active failover
  about 9-21
  actions 9-22
  configuring
    asymmetric routing support 9-45
  duplicate MAC addresses, avoiding 9-8
  primary status 9-21
  secondary status 9-21
Active/Standby failover
  about 9-19
  actions 9-20
  command replication 9-18
  configuration synchronization 9-18
  device initialization 9-18
  primary unit 9-19
  secondary unit 9-19
Adaptive Security Algorithm 1-22
add_acl 22-3
Add/Edit Access Group dialog box
  description 30-10
Add/Edit Filtering Entry dialog box
  description 27-27
Add/Edit IGMP Join Group dialog box
  description 30-8
Add/Edit IGMP Static Group dialog box
  description 30-10
Add/Edit Multicast Group dialog box 30-19
  description 30-19
Add/Edit OSPF Area dialog box 27-19

description 27-19
Add/Edit OSPF Neighbor Entry dialog box 27-24, 27-50, 27-51
description 27-25, 27-51
Add/Edit Rendezvous Point dialog box
restrictions 30-15
Add/Edit Summary Address dialog box
description 27-13, 27-19
Add/Edit Time Range dialog box 20-15
Add/Edit Virtual Link dialog box
description 27-28
admin context
about 8-2
changing 8-27
administrative access
using ICMP for 45-9
administrative distance 25-3, 25-4, 25-8
Advanced DHCP Options dialog box
description 17-6
Advanced OSPF Interface Properties dialog box 27-18
Advanced OSPF Virtual Link Properties dialog box
description 27-28
ae_standard_access_list_rule 22-3
ae_webtype_acl 23-3
AIP SSM
port-forwarding
enabling 13-7, 14-9
alternate address, ICMP message 48-15
analyzing syslog messages 41-2
application inspection
security level requirements 13-2, 14-2, 15-2
Apply button 4-12
Area/Networks tab
description 27-6
area border router 27-2
ARP inspection
about 6-6
enabling 6-11
static entry 6-10
ARP spoofing 6-6
ARP table
monitoring 11-40, 12-12, 13-23, 14-23
ARP test, failover 9-16
ASA (Adaptive Security Algorithm) 1-22
ASA 5505
Base license 12-2
MAC addresses 12-4
maximum VLANs 12-2
power over Ethernet 12-4
Security Plus license 12-2
SPAN 12-4
Spanning Tree Protocol, unsupported 12-8
ASA 5550 throughput 13-7, 14-10
ASA CX Staus tab 4-30
ASBR
definition of 27-2
ASR 9-45
ASR groups 9-45
asymmetric routing support 9-45
attributes
RADIUS 34-3
attribute-value pairs
TACACS+ 35-1
authenticating a certificate 40-12, 40-14
authentication
about 32-1
CLI access 45-17
Authentication tab
description 27-16
authorization
about 32-2
command 45-23
Auto-MDI/MDIX 11-2, 12-4
autostate messaging 2-13
Auto-Update, configuring 46-31

B

backed up configurations
  restoring 46-25
backing up configurations 46-21
Backing Up the Local CA Server 46-24
Baltimore Technologies, CA server support 40-5
bandwidth 4-20
bits subnet masks 48-3
BPDUs
  forwarding on the switch 2-13
broadcast Ping test 9-16
building blocks 20-1
bypassing the firewall, in the switch 2-6

C

CA
  CRs and 40-3
  public key cryptography 40-3
  revoked certificates 40-3
  supported servers 40-5
CA certificate 40-2
CA certificates 40-12, 40-14
Cancel button 4-12
CA server
  Digicert 40-5
  Geotrust 40-5
  Godaddy 40-5
  iPlanet 40-5
  Netscape 40-5
  RSA Keon 40-5
  Thawte 40-5
Catalyst 6500
  See switch
certificate
  CA 40-12, 40-14
  code-signer 40-29
  Identity 40-24
  local CA 40-31
certificate authentication 40-12, 40-14
certificate enrollment 40-13, 40-15
Certificate Revocation Lists
  See CRLs
change query interval 30-12
change query response time 30-12
change query timeout value 30-12
changing between contexts 8-25
changing the severity level 41-23
Cisco 7600
  See switch
Cisco IOS CS CA
  server support 40-5
Class A, B, and C addresses 48-1
classes, logging
  message class variables 41-4
  types 41-4
classes, resource
  See resource management
class map
  regular expression 20-14
Cluster Dashboard tab 4-25
Cluster Firewall Dashboard tab 4-26
clustering
  ASDM connection certificate IP address mismatch 10-12
  backup owner 10-10
  bootstrap configuration 10-50
  cabling 10-33
  cLACP
    system ID 10-62
    system priority 10-62
  cluster control link
    failure 10-9
    overview 10-7
    redundancy 10-8
    size 10-7
  configuration

examples 10-72
replication 10-11
configuring 10-59
connection
new, ownership 10-19
context mode 10-28
data path connection state replication 10-10
device-local EtherChannels, configuring on switch 10-30
executing a command cluster-wide 10-69
failover 10-28
feature history 10-88
features
centralized 10-21
individual units 10-22
NAT 10-24
SNMP 10-26
syslog and netflow 10-26
unsupported 10-20
VPN 10-26
guidelines and limitations 10-28
high availability 10-9
High Availability and Scalability Wizard 10-50
individual cluster interfaces, configuring 10-38
interface monitoring 10-9
IPv6 10-28
licensing 10-27
management
interface, configuring 10-38
interface, overview 10-11
network 10-11
overview 10-11
master unit
changing 10-67
election 10-3
maximum members 10-27
member requirements 10-3
model support 10-28
monitoring 10-70
overview
bootstrap configuration 10-3
cluster control link 10-7
Equal-Cost Multi-Path Routing 10-16
interfaces 10-4
load balancing 10-13
management 10-11
master unit 10-3
Policy-Based Routing 10-15
spanned EtherChannel 10-13
performance scaling factor 10-2
prerequisites 10-27
rebalancing new connections 10-20
removing a member 10-65
RSA key replication 10-12
software requirements 10-3
spanned EtherChannel
benefits 10-13
configuring 10-43
load balancing 10-14
maximum throughput 10-13
overview 10-13
redundancy 10-14
VSS or vPC 10-14
spanning-tree portfast 10-28
unit failure 10-9
unit health monitoring 10-9
upgrading software 10-3
code-signer certificate 40-29
command authorization
about 45-13
configuring 45-23
multiple contexts 45-14
configuration
factory default
commands 3-19
restoring 3-19
switch 2-1
configuration mode

accessing 3-2, 3-5
configurations, backing up 46-21
connection limits
per context 8-18
context mode 28-3
context modes 25-2, 26-3, 27-3, 29-3, 30-3
contexts
See security contexts
conversion error, ICMP message 48-16
creating a custom event list 41-18
CRL
cache refresh time 40-19, 40-23
custom messages list
logging output destination 41-5

D

data flow
routed firewall 6-13
transparent firewall 6-19
date and time in messages 41-22
default
class 8-9
routes, defining equal cost routes 25-7
default configuration
commands 3-19
restoring 3-19
default routes
about 25-7
configuring 25-7
device ID in messages 41-22
DHCP
monitoring
interface lease 13-24, 14-24
IP addresses 13-23, 14-23
server 13-23, 14-23
statistics 13-25, 14-25
relay 17-7
server 17-5
statistics 13-25, 14-25
DHCP Relay panel 18-7
DHCP services 16-5
digital certificates 40-1, 40-2
directory hierarchy search 36-3
disabling messages 41-22
DMZ, definition 1-18
DNS
server, configuring 16-8
dotted decimal subnet masks 48-3
dual IP stack, configuring 13-2
dual-ISP support 25-9
duplex
interface 12-9, 12-12
duplex, configuring 11-14, 12-6

E

echo reply, ICMP message 48-15
ECMP 25-3
Edit OSPF Interface Authentication dialog box 27-16
description 27-16
Edit OSPF Interface Properties dialog box 27-17
EIGRP
DUAL algorithm 28-2
hello interval 28-23
hello packets 28-1
hold time 28-2, 28-23
neighbor discovery 28-1
stub routing 28-6
stuck-in-active 28-2
enable command 3-2
enabling logging 41-7
enabling secure logging 41-21
enrolling
certificate 40-13, 40-15
Entrust, CA server support 40-5
established command, security level requirements 13-2, 14-2, 15-2

EtherChannel
adding interfaces 11-31
channel group 11-31
compatibility 11-5
converting existing interfaces 11-16
failover 11-12
guidelines 11-13
interface requirements 11-5
LACP 11-6
load balancing
configuring 11-33
overview 11-7
MAC address 11-8
management interface 11-30
maximum interfaces 11-33
minimum interfaces 11-33
mode
active 11-7
on 11-7
passive 11-7
overview 11-5
port priority 11-31
system priority 11-33
Ethernet
Auto-MDI/MDIX 11-2, 12-4
duplex 11-14, 12-6
jumbo frames, ASA 5580 11-39
jumbo frame support
single mode 13-14, 14-16
MTU 13-14, 14-16
speed 11-14, 12-6
evaluation license 5-24

F

factory default configuration
commands 3-19
restoring 3-19
failover
about 9-1
Active/Active, *See* Active/Active failover
Active/Standby, *See* Active/Standby failover
configuration file
terminal messages, Active/Standby 9-18
contexts 9-19
criteria 9-41
debug messages 9-53
disabling 9-51
enabling Stateful Failover 9-49
Ethernet failover cable 9-4
failover link 9-3
forcing 9-50
guidelines 42-5
health monitoring 9-16
interface health 9-16
interface monitoring 9-16
interface tests 9-16
link communications 9-3
MAC addresses
about 9-19
automatically assigning 8-12
module placement
inter-chassis 9-9
intra-chassis 9-8
monitoring, health 9-16
network tests 9-16
primary unit 9-19
redundant interfaces 11-12
reset 9-54
restoring a failed group 9-51
restoring a failed unit 9-51
secondary unit 9-19
SNMP syslog traps 9-53
Stateful Failover, *See* Stateful Failover
state link 9-4
switch configuration 2-13
system log messages 9-52
system requirements 9-2

trunk 2-13
unit health 9-16
failover groups
monitoring 9-54
reset 9-54
fast path 1-22
fiber interfaces 11-14
Fibre Channel interfaces
default settings 22-2, 23-3
filtering
security level requirements 13-2, 14-2, 15-2
filtering messages 41-4
editing 41-27
Filtering pane
description 27-27
firewall mode
about 6-1
configuring 6-1
flash memory available for logs 41-20
flow control for 10 Gigabit Ethernet 11-25
flow-export actions 43-4
format of messages 41-3
fragment protection 1-19

G

graphs
bookmarking 11-42, 12-15, 13-28, 14-28
interface monitoring 11-42, 12-15, 13-28, 14-28
printing 11-42, 12-15, 13-28, 14-28
groups
SNMP 42-3

H

H.323
transparent firewall guidelines 6-6
Help button 4-12
Help menu 4-9
high availability
about 9-1
history metrics 4-34
host
SNMP 42-3
hosts, subnet masks for 48-3
HSRP 6-5
HTTPS/Telnet/SSH
allowing network or host access to ASDM 45-1

I

ICMP
rules for access to ADSM 45-9
type numbers 48-15
ICMP unreachable message limits 45-10
Identity Certificates 40-24
implementing SNMP 42-4
individual syslog messages
assigning or changing rate limits 41-24
information reply, ICMP message 48-15
information request, ICMP message 48-15
inside, definition 1-18
installation
module verification 2-8
interface
duplex 12-9, 12-12
MTU 13-14, 14-16
status 4-20
subinterface, adding 11-38
throughput 4-20
Interface pane 27-16
interfaces
ASA 5505
enabled status 12-8
MAC addresses 12-4
maximum VLANs 12-2
switch port configuration 12-8

trunk ports 12-10
ASA 5550 throughput 13-7, 14-10
default settings 22-2, 23-3
duplex 11-14, 12-6
failover monitoring 9-16
fiber 11-14
jumbo frame support
single mode 13-14, 14-16
MAC addresses
automatically assigning 8-24
monitoring 11-40, 12-13, 13-26, 14-26
redundant 11-27
SFP 11-14
speed 11-14, 12-6
subinterfaces 11-36
turning off 13-22, 14-22
turning on 13-22, 14-22
IOS
upgrading 2-3
IP addresses
classes 48-1
management, transparent firewall 14-8
management, transparent firewall (8.3 and earlier) 15-4
private 48-2
subnet mask 48-4
IPv6
configuring alongside IPv4 13-2
default route 25-8
dual IP stack 13-2
duplicate address detection 31-3
neighbor discovery 31-1
router advertisement messages 31-3
static neighbors 31-5
static routes 25-8
IPv6 addresses
anycast 48-9
format 48-5
multicast 48-8
prefixes 48-10
required 48-10
types of 48-6
unicast 48-6
IPv6 prefixes 31-15
IPX 2-6

J

Java console 41-13
Join Group pane
description 30-8
jumbo frames, ASA 5580 11-39
jumbo frame support
single mode 13-14, 14-16

K

Kerberos
configuring 34-15, 36-7, 37-3
key pairs 40-25

L

LACP 11-6
Layer 2 firewall
See transparent firewall
Layer 2 forwarding table
See MAC address table
LDAP
attribute mapping 36-5
configuring 34-15, 36-7, 37-3
directory search 36-3
hierarchy example 36-2
SASL 36-2
licenses
activation key
entering 5-37

location 5-35
obtaining 5-36
ASA 5505 5-3
ASA 5510 5-4, 5-9
ASA 5520 5-5
ASA 5540 5-6
ASA 5550 5-7
ASA 5580 5-8, 5-17
ASA 5585-X 5-16
default 5-24
evaluation 5-24
failover 5-35
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
logging 41-5
licensing requirements for SNMP 42-4
link up/down test 9-16
local CA 40-31
Local CA User Database 40-35
local user database
adding a user 33-3
configuring 33-3
lockout recovery 45-31
logging
classes
filtering messages by 41-4
types 41-4
filtering
by message list 41-5
by severity level 41-1
output destinations
internal buffer 41-1, 41-8
Telnet or SSH session 41-8
queue
changing the size of 41-21
configuring 41-21
logging feature history 41-28
logging queue
configuring 41-21
login
banner, configuring 45-6
console 3-1
enable 3-2
global configuration mode 3-2
session 3-4
SSH 3-4, 45-5
Telnet 3-4
log viewers
executing certain commands 41-27
loops, avoiding 2-13

M

MAC address
redundant interfaces 11-5
MAC addresses
ASA 5505 12-4
automatically assigning 8-24
failover 9-19
security context classification 8-3
MAC address table
about 6-19

built-in-switch 6-7
MAC learning, disabling 6-13
monitoring 11-40, 12-12, 13-26, 14-26
resource management 8-18
static entry 6-12
MAC learning, disabling 6-13
management interfaces
default settings 22-2, 23-3
management IP address, transparent firewall 14-8
management IP address, transparent firewall (8.3 and earlier) 15-4
man-in-the-middle attack 6-6
mask
reply, ICMP message 48-15
request, ICMP message 48-15
Master Passphrase 16-5
menus 4-4
message filtering 41-4
message list
filtering by 41-5
messages, logging
classes
about 41-4
list of 41-4
component descriptions 41-3
filtering by message list 41-5
format of 41-3
severity levels 41-3
messages classes 41-4
messages in EMBLEM format 41-19, 41-20
metacharacters, regular expression 20-11
mgmt0 interfaces
default settings 22-2, 23-3
MIBs for SNMP 42-13
Microsoft Windows CA, supported 40-5
mobile redirection, ICMP message 48-16
mode
context 8-16
firewall 6-1
monitoring
ARP table 11-40, 12-12, 13-23, 14-23
DHCP
interface lease 13-24, 14-24
IP addresses 13-23, 14-23
server 13-23, 14-23
statistics 13-25, 14-25
failover 9-16
failover groups 9-54
history metrics 4-34
interfaces 11-40, 12-13, 13-26, 14-26
MAC address table 11-40, 12-12, 13-26, 14-26
OSPF 27-65
SNMP 42-1
monitoring logging 41-25
monitoring NSEL 43-7
monitoring switch traffic, ASA 5505 12-4
MRoute pane
description 30-6
MSFC
overview 2-2
SVIs 2-6
MTU 13-14, 14-16
multicast traffic 6-5
multiple context mode
logging 41-2
See security contexts
multiple SVIs 2-5

N

NAT
disabling proxy ARP for global addresses 24-11
neighbor reachable time 31-3
neighbor solicitation messages 31-2
neighrbor advertisement messages 31-2
NetFlow
overview 43-1
NetFlow event

matching to configured collectors 43-6
Network Activity test 9-16
No Payload Encryption 5-33
NSEL and syslog messages
redundant messages 43-2
NSEL feature history 43-8
NSEL licensing requirements 43-4
NT server
configuring 34-15, 36-7, 37-3

O

open ports 48-14
Options menu 4-5
OSPF
area authentication 27-19
area MD5 authentication 27-19
area parameters 27-19
authentication key 27-15
authentication support 27-2
configuring authentication 27-16
cost 27-15
dead interval 27-15
defining a static neighbor 27-24, 27-50, 27-51
defining interface properties 27-17
interaction with NAT 27-2
interface parameters 27-14
interface properties 27-16, 27-17
link-state advertisement 27-2
logging neighbor states 27-26
LSAs 27-2
MD5 authentication 27-15
monitoring 27-65
NSSA 27-20
packet pacing 27-66, 27-67
processes 27-2
redistributing routes 27-7
route calculation timers 27-25
route summarization 27-13
OSPF parameters
dead interval 27-18
hello interval 27-18
retransmit interval 27-18
transmit delay 27-18
output destination 41-6
output destinations 41-1, 41-8
e-mail address 41-1, 41-8
SNMP management station 41-1, 41-8
Telnet or SSH session 41-1, 41-8
outside, definition 1-18
oversubscribing resources 8-10

P

packet
classifier 8-3
packet flow
routed firewall 6-13
transparent firewall 6-19
parameter problem, ICMP message 48-15
pause frames for flow control 11-25
PIM
shortest path tree settings 30-18
PoE 12-4
port-forwarding
enabling 13-7, 14-9
ports
open on device 48-14
TCP and UDP 48-11
power over Ethernet 12-4
primary unit, failover 9-19
printing
graphs 11-42, 12-15, 13-28, 14-28
private networks 48-2
privileged EXEC mode
accessing 3-5
privileged EXEC mode, accessing 3-2
privileged mode

accessing 3-2
Process Instances tab
description 27-6
Product Authorization Key 5-36
Properties tab 27-17
description 27-17
fields 27-17
protocol numbers and literal values 48-11
Protocol pane (PIM)
description 30-14
proxy ARP, disabling 24-11
public key cryptography 40-3

Q

queue, logging
changing the size of 41-21

R

RADIUS
attributes 34-3
configuring a server 34-15, 36-7, 37-3
support 34-2
rapid link failure detection 2-13
rate limit 41-23
redirect, ICMP message 48-15
redundant interface
EtherChannel
converting existing interfaces 11-16
redundant interfaces
configuring 11-27
failover 11-12
MAC address 11-5
setting the active interface 11-30
Registration Authority description 40-3
regular expression 20-10
reloading
context 8-29
Request Filter pane
description 30-17
Reset button 4-12
resetting the services module 2-14
resource management
about 8-10
class 8-17
configuring 8-8
default class 8-9
oversubscribing 8-10
resource types 8-18
unlimited 8-11
restoring backups 46-25
revoked certificates 40-3
RFCs for SNMP 42-12
RIP
authentication 29-2
definition of 29-1
enabling 29-4
support for 29-2
RIP panel
limitations 29-3
RIP Version 2 Notes 29-3
routed mode
about 6-1
setting 6-1
route map
definition 26-1
route maps
defining 26-4
uses 26-1
router
advertisement, ICMP message 48-15
solicitation, ICMP message 48-15
router advertisement messages 31-3
router advertisement transmission interval 31-10
router lifetime value 31-11
routes

about default 25-7
configuring default routes 25-7
configuring IPv6 default 25-8
configuring IPv6 static 25-8
configuring static routes 25-3
Route Summarization tab
description 27-6
Route Tree pane 30-18
description 30-18
rules
ICMP 45-9
running configuration
copying 46-21

S

same security level communication
enabling 13-20, 14-21, 15-17
SDI
configuring 34-15, 36-7, 37-3
secondary unit, failover 9-19
Secure Copy
configure server 46-13
security appliance
connecting to 3-1
managing licenses 5-1
security contexts
about 8-1
adding 8-20
admin context
about 8-2
changing 8-27
cascading 8-6
changing between 8-25
classifier 8-3
command authorization 45-14
configuration
URL, changing 8-28
logging in 8-7
MAC addresses
automatically assigning 8-24
classifying using 8-3
managing 8-1, 8-26
monitoring 8-31
MSFC compatibility 2-3
multiple mode, enabling 8-16
nesting or cascading 8-7
reloading 8-29
removing 8-26
resource management 8-10
unsupported features 8-14
security level
about 13-1
security models for SNMP 42-3
segment size
maximum and minimum 13-13
maximum and minimum, overview 11-8
sending messages to an e-mail address 41-14
sending messages to a specified output destination 41-21
sending messages to a syslog server 41-9
sending messages to a Telnet or SSH session 41-18
sending messages to the console port 41-17
sending messages to the internal log buffer 41-12
session management path 1-22
severity levels, of system log messages
changing 41-1
filtering by 41-1
list of 41-3
severity levels, of system messages
definition 41-3
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
single mode
backing up configuration 8-16
configuration 8-16
enabling 8-16
restoring 8-16
Smart Call Home monitoring 44-9
SNMP
about 42-1
failover 42-5
management station 41-1, 41-8
prerequisites 42-4
SNMP configuration 42-6
SNMP groups 42-3
SNMP hosts 42-3
SNMP management station
adding 42-7
SNMP monitoring 42-11
SNMP terminology 42-2
SNMP users 42-3
SNMP Version 3 42-3, 42-9
SNMP Versions 1 and 2c 42-8
software
version 4-30
source quench, ICMP message 48-15
SPAN 12-4
Spanning Tree Protocol, unsupported 12-8
SPAN session 2-8
speed, configuring 11-14, 12-6
SSH
concurrent connections 45-2
login 45-5
username 45-5
startup configuration
copying 46-21
Startup Wizard
acessing 7-1
licensing requirements 7-1
Stateful Failover
about 9-13
enabling 9-49
state information 9-13
state link 9-4
stateful inspection 1-22
state information 9-13
state link 9-4
static ARP entry 6-10
static bridge entry 6-12
Static Group pane
description 30-9
static routes
configuring 25-3
deleting 25-6
status bar 4-11
stealth firewall
See transparent firewall
stuck-in-active 28-2
subinterface
adding 11-38
subinterfaces, adding 11-36
subnet masks
/bits 48-3
about 48-2
address range 48-4
determining 48-3
dotted decimal 48-3
number of hosts 48-3
subordinate certificate 40-2
Summary Address pane
description 27-11
SVIs
configuring 2-12
multiple 2-5
overview 2-5
switch
assigning VLANs to module 2-9
autostate messaging 2-13
BPDU forwarding 2-13

configuration 2-1
failover compatibility with transparent firewall 2-13
failover configuration 2-13
trunk for failover 2-13
verifying module installation 2-8
switched virtual interfaces
See SVIs
switch MAC address table 6-7
switch ports
access ports 12-8
SPAN 12-4
trunk ports 12-10
syslogd server program 41-6
syslog message filtering
using log viewers 41-25
syslog messages
analyzing 41-2
syslog messaging for SNMP 42-11
syslog server
designating more than one as output destination 41-6
system configuration 8-2
system log messages
classes 41-4
classes of 41-4
configuring in groups
by message list 41-5
by severity level 41-1
disabling logging of 41-1
filtering by message class 41-4
output destinations 41-1, 41-8
syslog message server 41-8
Telnet or SSH session 41-8
severity levels
about 41-3
changing the severity level of a message 41-1

T

TACACS+
command authorization, configuring 45-27
configuring a server 34-15, 36-7, 37-3
TCP
connection limits per context 8-18
maximum segment size 13-13
maximum segment size, overview 11-8
ports and literal values 48-11
TCP MSS
overview 11-8
Telnet
allowing management access 45-1
concurrent connections 45-2
login 45-4
temporary license 5-24
threat detection
scanning statistics
system performance 4-24
time exceeded, ICMP message 48-15
timestamp reply, ICMP message 48-15
timestamp request, ICMP message 48-15
Tools menu 4-6
traceroute, enabling 4-7
traffic flow
routed firewall 6-13
transparent firewall 6-19
transparent firewall
about 6-2
ARP inspection
about 6-6
enabling 6-11
static entry 6-10
data flow 6-19
guidelines 6-8
H.323 guidelines 6-6
HSRP 6-5
MAC learning, disabling 6-13
management IP address 14-8
management IP address (8.3 and earlier) 15-4
multicast traffic 6-5

static bridge entry 6-12
unsupported features 6-9
VRRP 6-5
trunk, 802.1Q 11-36
trunk ports 12-10
Trusted Flow Acceleration
modes 6-8, 7-1, 21-2
trustpoint 40-4

U

UDP
connection limits per context 8-18
connection state information 1-22
ports and literal values 48-11
unprivileged mode
accessing 3-4
unreachable, ICMP message 48-15
unreachable messages
required for MTU discovery 45-9
upgrading
IOS 2-3
URLs
context configuration, changing 8-28
user EXEC mode
accessing 3-2
username
adding 33-3
users
SNMP 42-3
using clustering 41-5, 43-3

V

VeriSign, configuring CAs example 40-6
version
IPS software 4-30
virtual firewalls
See security contexts
Virtual Link
description 27-28
virtual reassembly 1-19
VLANs 11-36
802.1Q trunk 11-36
ASA 5505
MAC addresses 12-4
maximum 12-2
assigning to FWSM 2-9
interfaces 2-9
subinterfaces 11-36
VPN
address range, subnets 48-4
VPN flex license 5-24
VRRP 6-5

W

WCCP 19-1
web caching 19-1
Window menu 4-9
Wizards menu 4-8

X

XOFF frames 11-25