



Cisco ASA Series General Operations CLI Configuration Guide

Software Version 9.1

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.

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 CLI Configuration Guide
Copyright © 2012-2014 Cisco Systems, Inc. All rights reserved.



About This Guide	i
Document Objectives	i
Related Documentation	i
Conventions	i
Obtaining Documentation and Submitting a Service Request	ii

PART 1

Getting Started with the ASA

CHAPTER 1

Introduction to the Cisco ASA	1-1
Hardware and Software Compatibility	1-1
VPN Compatibility	1-1
New Features	1-2
New Features in ASA 9.1(5)	1-2
New Features in ASA 9.1(4)	1-3
New Features in ASA 9.1(3)	1-5
New Features in ASA 9.1(2)	1-7
New Features in ASA 9.1(1)	1-13
Firewall Functional Overview	1-14
Security Policy Overview	1-14
Permitting or Denying Traffic with Access Lists	1-15
Applying NAT	1-15
Protecting from IP Fragments	1-15
Using AAA for Through Traffic	1-15
Applying HTTP, HTTPS, or FTP Filtering	1-15
Applying Application Inspection	1-15
Sending Traffic to a Module	1-15
Applying QoS Policies	1-16
Applying Connection Limits and TCP Normalization	1-16
Enabling Threat Detection	1-16
Enabling the Botnet Traffic Filter	1-16
Configuring Cisco Unified Communications	1-16
Firewall Mode Overview	1-17
Stateful Inspection Overview	1-17
VPN Functional Overview	1-18

Security Context Overview 1-19

ASA Clustering Overview 1-19

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 2-5

Verifying the Module Installation 2-6

Assigning VLANs to the ASA Services Module 2-7

Using the MSFC as a Directly Connected Router (SVIs) 2-10

Configuring the Switch for ASA Failover 2-11

Assigning VLANs to the Secondary ASA Services Module 2-11

Adding a Trunk Between a Primary Switch and Secondary Switch 2-11

Ensuring Compatibility with Transparent Firewall Mode 2-11

Enabling Autostate Messaging for Rapid Link Failure Detection 2-11

Resetting the ASA Services Module 2-12

Monitoring the ASA Services Module 2-12

Feature History for the Switch for Use with the ASA Services Module 2-15

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 3-5

Logging Out of a Telnet Session 3-6

Configuring ASDM Access for Appliances 3-6

Accessing ASDM Using the Factory Default Configuration 3-6

Customizing ASDM Access (ASA 5505) 3-7

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
Factory Default Configurations	3-18
Restoring the Factory Default Configuration	3-19
ASA 5505 Default Configuration	3-20
ASA 5505 Routed Mode Default Configuration	3-20
ASA 5505 Transparent Mode Sample Configuration	3-22
ASA 5510 and Higher Default Configuration	3-24
Working with the Configuration	3-24
Saving Configuration Changes	3-25
Saving Configuration Changes in Single Context Mode	3-25
Saving Configuration Changes in Multiple Context Mode	3-25
Copying the Startup Configuration to the Running Configuration	3-26
Viewing the Configuration	3-27
Clearing and Removing Configuration Settings	3-27
Creating Text Configuration Files Offline	3-28
Applying Configuration Changes to Connections	3-28
Reloading the ASA	3-29

CHAPTER 4**Managing Feature Licenses 4-1**

Supported Feature Licenses Per Model	4-1
Licenses Per Model	4-1
License Notes	4-18
VPN License and Feature Compatibility	4-23
Information About Feature Licenses	4-23
Preinstalled License	4-24
Permanent License	4-24
Time-Based Licenses	4-24
Time-Based License Activation Guidelines	4-24
How the Time-Based License Timer Works	4-25
How Permanent and Time-Based Licenses Combine	4-25
Stacking Time-Based Licenses	4-26
Time-Based License Expiration	4-26
Shared AnyConnect Premium Licenses	4-27
Information About the Shared Licensing Server and Participants	4-27
Communication Issues Between Participant and Server	4-28

- Information About the Shared Licensing Backup Server 4-28
- Failover and Shared Licenses 4-29
 - Maximum Number of Participants 4-29
- Failover or ASA Cluster Licenses 4-30
 - Failover License Requirements and Exceptions 4-30
 - ASA Cluster License Requirements and Exceptions 4-30
 - How Failover or ASA Cluster Licenses Combine 4-31
 - Loss of Communication Between Failover or ASA Cluster Units 4-32
 - Upgrading Failover Pairs 4-32
- No Payload Encryption Models 4-32
- Licenses FAQ 4-33
- Guidelines and Limitations 4-33
- Configuring Licenses 4-35
 - Obtaining an Activation Key 4-35
 - Activating or Deactivating Keys 4-36
 - Configuring a Shared License 4-37
 - Configuring the Shared Licensing Server 4-37
 - Configuring the Shared Licensing Backup Server (Optional) 4-39
 - Configuring the Shared Licensing Participant 4-39
- Monitoring Licenses 4-40
 - Viewing Your Current License 4-40
 - Monitoring the Shared License 4-49
- Feature History for Licensing 4-50

CHAPTER 5

Configuring the Transparent or Routed Firewall 5-1

- Information About the Firewall Mode 5-1
 - Information About Routed Firewall Mode 5-1
 - Information About Transparent Firewall Mode 5-2
 - Using the Transparent Firewall in Your Network 5-2
 - Bridge Groups 5-3
 - Management Interface (ASA 5510 and Higher) 5-4
 - Allowing Layer 3 Traffic 5-4
 - Allowed MAC Addresses 5-5
 - Passing Traffic Not Allowed in Routed Mode 5-5
 - Passing Traffic For Routed-Mode Features 5-5
 - BPDU Handling 5-5
 - MAC Address vs. Route Lookups 5-6
 - ARP Inspection 5-6
 - MAC Address Table 5-7

Licensing Requirements for the Firewall Mode	5-7
Default Settings	5-7
Guidelines and Limitations	5-8
Setting the Firewall Mode	5-9
Configuring ARP Inspection for the Transparent Firewall	5-10
Task Flow for Configuring ARP Inspection	5-10
Adding a Static ARP Entry	5-10
Enabling ARP Inspection	5-11
Customizing the MAC Address Table for the Transparent Firewall	5-12
Adding a Static MAC Address	5-12
Setting the MAC Address Timeout	5-12
Disabling MAC Address Learning	5-13
Monitoring the Transparent Firewall	5-13
Monitoring ARP Inspection	5-13
Monitoring the MAC Address Table	5-13
Firewall Mode Examples	5-14
How Data Moves Through the ASA in Routed Firewall Mode	5-14
An Inside User Visits a Web Server	5-15
An Outside User Visits a Web Server on the DMZ	5-16
An Inside User Visits a Web Server on the DMZ	5-17
An Outside User Attempts to Access an Inside Host	5-17
A DMZ User Attempts to Access an Inside Host	5-19
How Data Moves Through the Transparent Firewall	5-20
An Inside User Visits a Web Server	5-21
An Inside User Visits a Web Server Using NAT	5-22
An Outside User Visits a Web Server on the Inside Network	5-23
An Outside User Attempts to Access an Inside Host	5-24
Feature History for the Firewall Mode	5-25

PART 2**Configuring High Availability and Scalability****CHAPTER 6****Configuring Multiple Context Mode 6-1**

Information About Security Contexts	6-1
Common Uses for Security Contexts	6-2
Context Configuration Files	6-2
Context Configurations	6-2
System Configuration	6-2
Admin Context Configuration	6-2

How the ASA Classifies Packets	6-3
Valid Classifier Criteria	6-3
Classification Examples	6-4
Cascading Security Contexts	6-6
Management Access to Security Contexts	6-7
System Administrator Access	6-7
Context Administrator Access	6-8
Information About Resource Management	6-8
Resource Classes	6-8
Resource Limits	6-8
Default Class	6-9
Using Oversubscribed Resources	6-10
Using Unlimited Resources	6-11
Information About MAC Addresses	6-11
Default MAC Address	6-12
Interaction with Manual MAC Addresses	6-12
Failover MAC Addresses	6-12
MAC Address Format	6-12
Licensing Requirements for Multiple Context Mode	6-13
Prerequisites	6-14
Guidelines and Limitations	6-14
Default Settings	6-15
Configuring Multiple Contexts	6-15
Task Flow for Configuring Multiple Context Mode	6-15
Enabling or Disabling Multiple Context Mode	6-16
Enabling Multiple Context Mode	6-16
Restoring Single Context Mode	6-16
Configuring a Class for Resource Management	6-17
Configuring a Security Context	6-19
Automatically Assigning MAC Addresses to Context Interfaces	6-24
Changing Between Contexts and the System Execution Space	6-24
Managing Security Contexts	6-25
Removing a Security Context	6-25
Changing the Admin Context	6-26
Changing the Security Context URL	6-26
Reloading a Security Context	6-27
Reloading by Clearing the Configuration	6-28
Reloading by Removing and Re-adding the Context	6-28
Monitoring Security Contexts	6-28

Viewing Context Information	6-29
Viewing Resource Allocation	6-30
Viewing Resource Usage	6-33
Monitoring SYN Attacks in Contexts	6-34
Viewing Assigned MAC Addresses	6-36
Viewing MAC Addresses in the System Configuration	6-37
Viewing MAC Addresses Within a Context	6-38
Configuration Examples for Multiple Context Mode	6-39
Feature History for Multiple Context Mode	6-40

CHAPTER 7**Configuring Failover 7-1**

Introduction to Failover	7-1
Failover Overview	7-2
Failover System Requirements	7-2
Hardware Requirements	7-2
Software Requirements	7-2
License Requirements	7-3
Failover and Stateful Failover Links	7-3
Failover Link	7-3
Stateful Failover Link	7-4
Avoiding Interrupted Failover and Data Links	7-5
MAC Addresses and IP Addresses	7-7
Intra- and Inter-Chassis Module Placement for the ASA Services Module	7-8
Intra-Chassis Failover	7-8
Inter-Chassis Failover	7-9
Stateless and Stateful Failover	7-12
Stateless Failover	7-13
Stateful Failover	7-13
Transparent Firewall Mode Requirements	7-14
Transparent Mode Requirements for Appliances	7-15
Transparent Mode Requirements for Modules	7-15
Failover Health Monitoring	7-16
Unit Health Monitoring	7-16
Interface Monitoring	7-17
Failover Times	7-18
Configuration Synchronization	7-18
Running Configuration Replication	7-18
Command Replication	7-19
Information About Active/Standby Failover	7-20

Primary/Secondary Roles and Active/Standby Status	7-20
Active Unit Determination at Startup	7-20
Failover Events	7-20
Information About Active/Active Failover	7-21
Active/Active Failover Overview	7-22
Primary/Secondary Roles and Active/Standby Status for a Failover Group	7-22
Active Unit Determination for Failover Groups at Startup	7-22
Failover Events	7-23
Licensing Requirements Failover	7-24
Prerequisites for Failover	7-24
Guidelines and Limitations	7-25
Default Settings	7-25
Configuring Active/Standby Failover	7-26
Configuring the Primary Unit for Active/Standby Failover	7-26
Configuring the Secondary Unit for Active/Standby Failover	7-30
Configuring Active/Active Failover	7-30
Configuring the Primary Unit for Active/Active Failover	7-31
Configuring the Secondary Unit for Active/Active Failover	7-35
Configuring Optional Failover Parameters	7-35
Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses	7-36
Configuring Interface Monitoring	7-38
Configuring Support for Asymmetrically Routed Packets (Active/Active Mode)	7-39
Managing Failover	7-42
Forcing Failover	7-42
Disabling Failover	7-43
Restoring a Failed Unit	7-44
Re-Syncing the Configuration	7-44
Testing the Failover Functionality	7-44
Remote Command Execution	7-45
Sending a Command	7-45
Changing Command Modes	7-46
Security Considerations	7-47
Limitations of Remote Command Execution	7-47
Monitoring Failover	7-48
Failover Messages	7-48
Failover Syslog Messages	7-48
Failover Debug Messages	7-48
SNMP Failover Traps	7-48
Monitoring Failover	7-49

Feature History for Failover 7-49

CHAPTER 8

Configuring a Cluster of ASAs 8-1

- Information About ASA Clustering 8-1
 - How the ASA Cluster Fits into Your Network 8-2
 - Performance Scaling Factor 8-2
 - Cluster Members 8-2
 - ASA Hardware and Software Requirements 8-3
 - Bootstrap Configuration 8-3
 - Master and Slave Unit Roles 8-3
 - Master Unit Election 8-3
 - ASA Cluster Interfaces 8-4
 - Interface Types 8-4
 - Interface Type Mode 8-6
 - Cluster Control Link 8-6
 - Cluster Control Link Traffic Overview 8-7
 - Cluster Control Link Interfaces and Network 8-7
 - Sizing the Cluster Control Link 8-7
 - Cluster Control Link Redundancy 8-8
 - Cluster Control Link Reliability 8-8
 - Cluster Control Link Failure 8-9
 - High Availability within the ASA Cluster 8-9
 - Unit Health Monitoring 8-9
 - Interface monitoring 8-9
 - Unit or Interface Failure 8-9
 - Rejoining the Cluster 8-10
 - Data Path Connection State Replication 8-10
 - Configuration Replication 8-11
 - ASA Cluster Management 8-11
 - Management Network 8-11
 - Management Interface 8-11
 - Master Unit Management Vs. Slave Unit Management 8-12
 - RSA Key Replication 8-12
 - ASDM Connection Certificate IP Address Mismatch 8-12
 - Load Balancing Methods 8-13
 - Spanned EtherChannel (Recommended) 8-13
 - Policy-Based Routing (Routed Firewall Mode Only) 8-15
 - Equal-Cost Multi-Path Routing (Routed Firewall Mode Only) 8-16
 - Inter-Site Clustering 8-16
 - Inter-Site Clustering Guidelines 8-16

Sizing the Data Center Interconnect	8-17
Inter-Site Example	8-17
How the ASA Cluster Manages Connections	8-18
Connection Roles	8-18
New Connection Ownership	8-19
Sample Data Flow	8-19
Rebalancing New TCP Connections Across the Cluster	8-20
ASA Features and Clustering	8-20
Unsupported Features	8-20
Centralized Features	8-21
Features Applied to Individual Units	8-22
Dynamic Routing	8-22
Multicast Routing	8-24
NAT	8-24
AAA for Network Access	8-25
Syslog and Netflow	8-26
SNMP	8-26
VPN	8-26
FTP	8-26
Cisco TrustSec	8-27
Licensing Requirements for ASA Clustering	8-27
Prerequisites for ASA Clustering	8-27
Guidelines and Limitations	8-28
Default Settings	8-32
Configuring ASA Clustering	8-32
Task Flow for ASA Cluster Configuration	8-32
Cabling the Cluster Units and Configuring Upstream and Downstream Equipment	8-33
Configuring the Cluster Interface Mode on Each Unit	8-35
Configuring Interfaces on the Master Unit	8-36
Configuring Individual Interfaces (Recommended for the Management Interface)	8-36
Configuring Spanned EtherChannels	8-38
Configuring the Master Unit Bootstrap Settings	8-42
Prerequisites	8-42
Enabling the Cluster Control Link Interface	8-43
Configuring Basic Bootstrap Settings and Enabling Clustering	8-45
Configuring Advanced Clustering Settings	8-47
Examples	8-48
Configuring Slave Unit Bootstrap Settings	8-48
Prerequisites	8-49

Enabling the Cluster Control Link Interface	8-49
Configuring Bootstrap Settings and Joining the Cluster	8-50
Examples	8-52
Managing ASA Cluster Members	8-53
Becoming an Inactive Member	8-54
Inactivating a Member	8-55
Leaving the Cluster	8-56
Changing the Master Unit	8-57
Executing a Command Cluster-Wide	8-58
Monitoring the ASA Cluster	8-59
Monitoring Commands	8-59
Related Commands	8-61
Configuration Examples for ASA Clustering	8-63
Sample ASA and Switch Configuration	8-63
ASA Configuration	8-63
IOS Switch Configuration	8-65
Firewall on a Stick	8-66
Traffic Segregation	8-68
Redundant Interface (PBR or ECMP)	8-70
Spanned EtherChannel With Backup Links	8-72
Feature History for ASA Clustering	8-78

PART 3**Configuring Interfaces****CHAPTER 9****Starting Interface Configuration (ASA 5510 and Higher) 9-1**

Information About Starting ASA 5510 and Higher Interface Configuration	9-2
Auto-MDI/MDIX Feature	9-2
Interfaces in Transparent Mode	9-2
Management Interface	9-2
Management Interface Overview	9-2
Management <i>Slot/Port</i> Interface	9-3
Using Any Interface for Management-Only Traffic	9-3
Management Interface for Transparent Mode	9-4
No Support for Redundant Management Interfaces	9-4
Management O/O Interface on the ASA 5512-X through ASA 5555-X	9-4
Redundant Interfaces	9-5
Redundant Interface MAC Address	9-5
EtherChannels	9-5
Channel Group Interfaces	9-5

Connecting to an EtherChannel on Another Device	9-5
Link Aggregation Control Protocol	9-6
Load Balancing	9-7
EtherChannel MAC Address	9-8
Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size	9-8
MTU Overview	9-8
Default MTU	9-8
Path MTU Discovery	9-9
Setting the MTU and Jumbo Frames	9-9
TCP Maximum Segment Size Overview	9-9
Default TCP MSS	9-9
Setting the TCP MSS for VPN and Non-VPN Traffic	9-9
Examples	9-10
Licensing Requirements for ASA 5510 and Higher Interfaces	9-10
Guidelines and Limitations	9-12
Default Settings	9-14
Starting Interface Configuration (ASA 5510 and Higher)	9-15
Task Flow for Starting Interface Configuration	9-15
Converting In-Use Interfaces to a Redundant or EtherChannel Interface	9-16
Enabling the Physical Interface and Configuring Ethernet Parameters	9-26
Configuring a Redundant Interface	9-28
Configuring a Redundant Interface	9-28
Changing the Active Interface	9-30
Configuring an EtherChannel	9-30
Adding Interfaces to the EtherChannel	9-30
Customizing the EtherChannel	9-32
Configuring VLAN Subinterfaces and 802.1Q Trunking	9-33
Enabling Jumbo Frame Support (Supported Models)	9-35
Monitoring Interfaces	9-36
Configuration Examples for ASA 5510 and Higher Interfaces	9-36
Physical Interface Parameters Example	9-36
Subinterface Parameters Example	9-37
Multiple Context Mode Example	9-37
EtherChannel Example	9-37
Where to Go Next	9-37
Feature History for ASA 5510 and Higher Interfaces	9-38

CHAPTER 10**Starting Interface Configuration (ASA 5505) 10-1**

- Information About ASA 5505 Interfaces 10-1
 - Understanding ASA 5505 Ports and Interfaces 10-2
 - Maximum Active VLAN Interfaces for Your License 10-2
 - VLAN MAC Addresses 10-4
 - Power over Ethernet 10-4
 - Monitoring Traffic Using SPAN 10-4
 - Auto-MDI/MDIX Feature 10-4
- Licensing Requirements for ASA 5505 Interfaces 10-4
- Guidelines and Limitations 10-5
- Default Settings 10-5
- Starting ASA 5505 Interface Configuration 10-6
 - Task Flow for Starting Interface Configuration 10-6
 - Configuring VLAN Interfaces 10-6
 - Configuring and Enabling Switch Ports as Access Ports 10-7
 - Configuring and Enabling Switch Ports as Trunk Ports 10-9
- Monitoring Interfaces 10-11
- Configuration Examples for ASA 5505 Interfaces 10-11
 - Access Port Example 10-11
 - Trunk Port Example 10-12
- Where to Go Next 10-13
- Feature History for ASA 5505 Interfaces 10-13

CHAPTER 11**Completing Interface Configuration (Routed Mode) 11-1**

- Information About Completing Interface Configuration in Routed Mode 11-1
 - Security Levels 11-1
 - Dual IP Stack (IPv4 and IPv6) 11-2
- Licensing Requirements for Completing Interface Configuration in Routed Mode 11-2
- Guidelines and Limitations 11-5
- Default Settings 11-6
- Completing Interface Configuration in Routed Mode 11-6
 - Task Flow for Completing Interface Configuration 11-7
 - Configuring General Interface Parameters 11-7
 - Configuring the MAC Address, MTU, and TCP MSS 11-10
 - Configuring IPv6 Addressing 11-12
 - Information About IPv6 11-12
 - Configuring a Global IPv6 Address 11-13
 - Configuring IPv6 Neighbor Discovery 11-15

- Allowing Same Security Level Communication 11-15
- Turning Off and Turning On Interfaces 11-17
- Monitoring Interfaces 11-17
- Configuration Examples for Interfaces in Routed Mode 11-18
 - ASA 5505 Example 11-18
- Feature History for Interfaces in Routed Mode 11-19

CHAPTER 12

- Completing Interface Configuration (Transparent Mode) 12-1**
 - Information About Completing Interface Configuration in Transparent Mode 12-1
 - Bridge Groups in Transparent Mode 12-2
 - Security Levels 12-2
 - Licensing Requirements for Completing Interface Configuration in Transparent Mode 12-3
 - Guidelines and Limitations 12-5
 - Default Settings 12-7
 - Completing Interface Configuration in Transparent Mode 12-7
 - Task Flow for Completing Interface Configuration 12-8
 - Configuring Bridge Groups 12-8
 - Configuring General Interface Parameters 12-9
 - Configuring a Management Interface (ASA 5510 and Higher) 12-12
 - Configuring the MAC Address, MTU, and TCP MSS 12-13
 - Configuring IPv6 Addressing 12-16
 - Information About IPv6 12-16
 - Configuring a Global IPv6 Address 12-17
 - Configuring IPv6 Neighbor Discovery 12-18
 - Allowing Same Security Level Communication 12-18
 - Turning Off and Turning On Interfaces 12-19
 - Monitoring Interfaces 12-19
 - Configuration Examples for Interfaces in Transparent Mode 12-20
 - Feature History for Interfaces in Transparent Mode 12-21

PART 4

Configuring Basic Settings

CHAPTER 13

- Configuring Basic Settings 13-1**
 - Configuring the Hostname, Domain Name, and Passwords 13-1
 - Setting the Login Password 13-2
 - Changing the Enable Password 13-3
 - Setting the Hostname 13-3
 - Setting the Domain Name 13-4

Feature History for the Hostname, Domain Name, and Passwords	13-4
Setting the Date and Time	13-4
Setting the Time Zone and Daylight Saving Time Date Range	13-5
Setting the Date and Time Using an NTP Server	13-6
Setting the Date and Time Manually	13-8
Configuring the Master Passphrase	13-8
Information About the Master Passphrase	13-8
Licensing Requirements for the Master Passphrase	13-9
Guidelines and Limitations	13-9
Adding or Changing the Master Passphrase	13-9
Disabling the Master Passphrase	13-11
Recovering the Master Passphrase	13-12
Feature History for the Master Passphrase	13-13
Configuring the DNS Server	13-13
Performing Password Recovery	13-14
Recovering Passwords for the ASA	13-14
Disabling Password Recovery	13-16
Monitoring DNS Cache	13-16

CHAPTER 14**Configuring DHCP Services 14-1**

Information About DHCP Services	14-1
Information About the DHCP Server	14-1
Information About the DHCP Relay Agent	14-2
Licensing Requirements for DHCP	14-2
Guidelines and Limitations	14-2
Configuring DHCP Services	14-4
Configuring the DHCP Server	14-4
Enabling the DHCP Server	14-5
Configuring DHCP Options	14-6
Configuring the DHCP Relay Agent	14-8
Configuring the DHCPv4 Relay Agent	14-8
Configuring the DHCPv6 Relay Agent	14-10
Additional References	14-11
RFCs	14-11
Monitoring DHCP Services	14-11
Feature History for DHCP Services	14-12

CHAPTER 15

Configuring Dynamic DNS 15-1

Information About DDNS 15-1

Licensing Requirements for DDNS 15-2

Guidelines and Limitations 15-2

Configuring DDNS 15-2

Configuration Examples for DDNS 15-3

Example 1: Client Updates Both A and PTR RRs for Static IP Addresses 15-3

Example 2: Client Updates Both A and PTR RRs; DHCP Server Honors Client Update Request; FQDN Provided Through Configuration 15-4

Example 3: Client Includes FQDN Option Instructing Server Not to Update Either RR; Server Overrides Client and Updates Both RRs. 15-5

Example 4: Client Asks Server To Perform Both Updates; Server Configured to Update PTR RR Only; Honors Client Request and Updates Both A and PTR RR 15-6

Example 5: Client Updates A RR; Server Updates PTR RR 15-7

DDNS Monitoring Commands 15-9

Feature History for DDNS 15-9

CHAPTER 16

Configuring Web Cache Services Using WCCP 16-1

Information About WCCP 16-1

Guidelines and Limitations 16-1

Licensing Requirements for WCCP 16-2

Enabling WCCP Redirection 16-3

WCCP Monitoring Commands 16-4

Feature History for WCCP 16-4

PART 5

Configuring Objects and ACLs

CHAPTER 17

Configuring Objects 17-1

Information About Objects 17-1

Licensing Requirements for Objects 17-1

Guidelines and Limitations 17-1

Configuring Objects 17-2

Configuring Network Objects and Groups 17-2

Configuring a Network Object 17-2

Configuring a Network Object Group 17-3

Configuring Service Objects and Service Groups 17-5

Configuring a Service Object 17-5

Configuring a Service Group 17-6

Configuring a TCP or UDP Port Service Group	17-8
Configuring an ICMP Group	17-10
Configuring a Protocol Group	17-11
Configuring Local User Groups	17-11
Configuring Security Group Object Groups	17-13
Configuring Regular Expressions	17-14
Creating a Regular Expression	17-14
Creating a Regular Expression Class Map	17-17
Configuring Time Ranges	17-18
Monitoring Objects	17-19
Feature History for Objects	17-19

CHAPTER 18**Information About Access Control Lists 18-1**

ACL Types	18-1
Access Control Entry Order	18-2
Access Control Implicit Deny	18-3
IP Addresses Used for ACLs When You Use NAT	18-3
Where to Go Next	18-3

CHAPTER 19**Adding an Extended Access Control List 19-1**

Information About Extended ACLs	19-1
Access Control Entry Order	19-1
NAT and ACLs	19-2
Information About Scheduling ACL Activation	19-2
Licensing Requirements for Extended ACLs	19-3
Guidelines and Limitations	19-3
Default Settings	19-4
Configuring Extended ACLs	19-4
Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy	19-4
Adding an ACE for TCP or UDP-Based Policy, with Ports	19-6
Adding an ACE for ICMP-Based Policy, with ICMP Type	19-7
Adding an ACE for User-Based Policy (Identity Firewall)	19-7
Adding an ACE for Security Group-Based Policy (TrustSec)	19-8
Adding Remarks to ACLs	19-9
Monitoring Extended ACLs	19-10
Configuration Examples for Extended ACLs	19-10
Configuration Examples for Extended ACLs (No Objects)	19-10
Configuration Examples for Extended ACLs (Using Objects)	19-11

Where to Go Next 19-12
 Feature History for Extended ACLs 19-12

CHAPTER 20

Adding an EtherType Access Control List 20-1
 Information About EtherType ACLs 20-1
 Licensing Requirements for EtherType ACLs 20-1
 Guidelines and Limitations 20-2
 Default Settings 20-2
 Configuring EtherType ACLs 20-2
 Task Flow for Configuring EtherType ACLs 20-2
 Adding EtherType ACLs 20-3
 Adding Remarks to ACLs 20-4
 What to Do Next 20-4
 Monitoring EtherType ACLs 20-4
 Configuration Examples for EtherType ACLs 20-5
 Feature History for EtherType ACLs 20-5

CHAPTER 21

Adding a Standard Access Control List 21-1
 Information About Standard ACLs 21-1
 Licensing Requirements for Standard ACLs 21-1
 Guidelines and Limitations 21-1
 Default Settings 21-2
 Adding Standard ACLs 21-3
 Task Flow for Configuring Extended ACLs 21-3
 Adding a Standard ACL 21-3
 Adding Remarks to ACLs 21-4
 What to Do Next 21-4
 Monitoring ACLs 21-4
 Configuration Examples for Standard ACLs 21-4
 Feature History for Standard ACLs 21-5

CHAPTER 22

Adding a Webtype Access Control List 22-1
 Licensing Requirements for Webtype ACLs 22-1
 Guidelines and Limitations 22-1
 Default Settings 22-3
 Using Webtype ACLs 22-3
 Task Flow for Configuring Webtype ACLs 22-3

Adding Webtype ACLs with a URL String	22-4
Adding Webtype ACLs with an IP Address	22-5
Adding Remarks to ACLs	22-5
What to Do Next	22-6
Monitoring Webtype ACLs	22-6
Configuration Examples for Webtype ACLs	22-6
Feature History for Webtype ACLs	22-8

CHAPTER 23**Configuring Logging for Access Control Lists 23-1**

Configuring Logging for ACLs	23-1
Information About Logging ACL Activity	23-1
Licensing Requirements for ACL Logging	23-2
Guidelines and Limitations	23-2
Default Settings	23-3
Configuring ACL Logging	23-3
Monitoring ACLs	23-4
Configuration Examples for ACL Logging	23-4
Feature History for ACL Logging	23-5
Managing Deny Flows	23-5
Information About Managing Deny Flows	23-6
Licensing Requirements for Managing Deny Flows	23-6
Guidelines and Limitations	23-6
Default Settings	23-7
Managing Deny Flows	23-7
Monitoring Deny Flows	23-7
Feature History for Managing Deny Flows	23-8

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 a Default Static Route 25-4
 - Limitations on Configuring a Default Static Route 25-4
 - Configuring IPv6 Default and Static Routes 25-5
 - Monitoring a Static or Default Route 25-6
 - Configuration Examples for Static or Default Routes 25-8
 - Feature History for Static and Default Routes 25-9

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
 - Customizing a Route Map 26-4

CHAPTER 27**Configuring OSPF 27-1**

- Defining a Route to Match a Specific Destination Address 26-4
 - Configuring the Metric Values for a Route Action 26-5
 - Configuration Example for Route Maps 26-6
 - Feature History for Route Maps 26-6
- 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
- Configuring EIGRP 28-3
 - Enabling EIGRP 28-4
 - Enabling EIGRP Stub Routing 28-4
- Customizing EIGRP 28-5
 - Defining a Network for an EIGRP Routing Process 28-6
 - Configuring Interfaces for EIGRP 28-7
 - Configuring Passive Interfaces 28-8
 - Configuring the Summary Aggregate Addresses on Interfaces 28-9
 - Changing the Interface Delay Value 28-10
 - Enabling EIGRP Authentication on an Interface 28-10
 - Defining an EIGRP Neighbor 28-12
 - Redistributing Routes Into EIGRP 28-12
 - Filtering Networks in EIGRP 28-14
 - Customizing the EIGRP Hello Interval and Hold Time 28-15
 - Disabling Automatic Route Summarization 28-16
 - Configuring Default Information in EIGRP 28-16
 - Disabling EIGRP Split Horizon 28-17
 - Restarting the EIGRP Process 28-18
- Monitoring EIGRP 28-18
- Configuration Example for EIGRP 28-19

Feature History for EIGRP 28-20

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 29-4

Configuring the RIP Version 29-5

Configuring Interfaces for RIP 29-6

Configuring the RIP Send and Receive Version on an Interface 29-6

Configuring Route Summarization 29-7

Filtering Networks in RIP 29-8

Redistributing Routes into the RIP Routing Process 29-8

Enabling RIP Authentication 29-9

Restarting the RIP Process 29-10

Monitoring RIP 29-11

Configuration Example for RIP 29-11

Feature History for RIP 29-12

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 30-2

Clustering 30-2

Licensing Requirements for Multicast Routing 30-3

Guidelines and Limitations 30-3

Enabling Multicast Routing 30-3

Customizing Multicast Routing 30-4

Configuring Stub Multicast Routing and Forwarding IGMP Messages 30-4

- Configuring a Static Multicast Route 30-5
- Configuring IGMP Features 30-5
 - Disabling IGMP on an Interface 30-6
 - Configuring IGMP Group Membership 30-7
 - Configuring a Statically Joined IGMP Group 30-7
 - Controlling Access to Multicast Groups 30-8
 - Limiting the Number of IGMP States on an Interface 30-8
 - Modifying the Query Messages to Multicast Groups 30-8
 - Changing the IGMP Version 30-9
- Configuring PIM Features 30-10
 - Enabling and Disabling PIM on an Interface 30-10
 - Configuring a Static Rendezvous Point Address 30-11
 - Configuring the Designated Router Priority 30-11
 - Configuring and Filtering PIM Register Messages 30-12
 - Configuring PIM Message Intervals 30-12
 - Filtering PIM Neighbors 30-12
- Configuring a Bidirectional Neighbor Filter 30-13
- Configuring a Multicast Boundary 30-14
- Configuration Example for Multicast Routing 30-15
- Additional References 30-15
 - Related Documents 30-16
 - RFCs 30-16
- Feature History for Multicast Routing 30-16

CHAPTER 31

- Configuring IPv6 Neighbor Discovery 31-1**
 - Information About IPv6 Neighbor Discovery 31-1
 - Neighbor Solicitation Messages 31-2
 - Neighbor Reachable Time 31-2
 - Duplicate Address Detection 31-2
 - Router Advertisement Messages 31-3
 - Static IPv6 Neighbors 31-4
 - Licensing Requirements for IPv6 Neighbor Discovery 31-4
 - Prerequisites for IPv6 Neighbor Discovery 31-4
 - Guidelines and Limitations 31-4
 - Default Settings for IPv6 Neighbor Discovery 31-6
 - Configuring IPv6 Neighbor Discovery 31-6
 - Entering Interface Configuration Mode 31-6
 - Configuring the Neighbor Solicitation Message Interval 31-7
 - Configuring the Neighbor Reachable Time 31-8

Configuring the Router Advertisement Transmission Interval	31-8
Configuring the Router Lifetime Value	31-9
Configuring DAD Settings	31-9
Suppressing Router Advertisement Messages	31-10
Configuring Address Config Flags for IPv6 DHCP Relay	31-11
Configuring the IPv6 Prefix in Router Advertisements	31-12
Configuring a Static IPv6 Neighbor	31-13
Monitoring IPv6 Neighbor Discovery	31-14
Additional References	31-14
Related Documents for IPv6 Prefixes	31-15
RFCs for IPv6 Prefixes and Documentation	31-15
Feature History for IPv6 Neighbor Discovery	31-15

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	33-2
Guidelines and Limitations	33-3
Adding a User Account to the Local Database	33-4
Monitoring the Local Database	33-8
Feature History for the Local Database	33-9

CHAPTER 34**Configuring RADIUS Servers for AAA** 34-1

Information About RADIUS Servers	34-1
----------------------------------	------

- Supported Authentication Methods 34-1
- 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 34-17
- 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 35-5
- Monitoring TACACS+ Servers 35-6
- 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
Configuring Authorization with LDAP for VPN	36-10
Monitoring LDAP Servers	36-11
Feature History for LDAP Servers	36-12

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	37-4
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 Agents	38-13
Configuring Identity Options	38-14
Configuring Identity-Based Security Policy	38-19
Collecting User Statistics	38-20
Configuration Examples	38-21
AAA Rule and Access Rule Example 1	38-21
AAA Rule and Access Rule Example 2	38-21

- VPN Filter Example 38-22
 - VPN with IDFW Rule -1 Example 38-22
 - VPN with IDFW Rule -2 Example 38-22
- Monitoring the Identity Firewall 38-23
 - Monitoring AD Agents 38-23
 - Monitoring Groups 38-23
 - Monitoring Memory Usage for the Identity Firewall 38-23
 - Monitoring Users for the Identity Firewall 38-24
- Feature History for the Identity Firewall 38-25

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 39-17
 - Configuring the Security Exchange Protocol (SXP) 39-19
 - Adding an SXP Connection Peer 39-22
 - Refreshing Environment Data 39-23
 - Configuring the Security Policy 39-23
- Configuration Example 39-25
- Monitoring Cisco TrustSec 39-25

Feature History for the Cisco TrustSec Integration 39-26

CHAPTER 40

Configuring Digital Certificates	40-1
Information About Digital Certificates	40-1
Public Key Cryptography	40-2
Certificate Scalability	40-2
Key Pairs	40-2
Trustpoints	40-3
Certificate Enrollment	40-3
Proxy for SCEP Requests	40-3
Revocation Checking	40-4
Supported CA Servers	40-4
CRLs	40-5
OCSP	40-5
The Local CA	40-6
Storage for Local CA Files	40-6
The Local CA Server	40-7
Using Certificates and User Login Credentials	40-7
Using User Login Credentials	40-7
Using Certificates	40-8
Licensing Requirements for Digital Certificates	40-8
Prerequisites for Local Certificates	40-9
Prerequisites for SCEP Proxy Support	40-9
Guidelines and Limitations	40-9
Configuring Digital Certificates	40-10
Configuring Key Pairs	40-11
Removing Key Pairs	40-12
Configuring Trustpoints	40-12
Configuring CRLs for a Trustpoint	40-15
Exporting a Trustpoint Configuration	40-17
Importing a Trustpoint Configuration	40-18
Configuring CA Certificate Map Rules	40-19
Obtaining Certificates Manually	40-20
Obtaining Certificates Automatically with SCEP	40-22
Configuring Proxy Support for SCEP Requests	40-23
Enabling the Local CA Server	40-24
Configuring the Local CA Server	40-25
Customizing the Local CA Server	40-27
Debugging the Local CA Server	40-28

- Disabling the Local CA Server 40-28
- Deleting the Local CA Server 40-28
- Configuring Local CA Certificate Characteristics 40-29
 - Configuring the Issuer Name 40-30
 - Configuring the CA Certificate Lifetime 40-30
 - Configuring the User Certificate Lifetime 40-31
 - Configuring the CRL Lifetime 40-32
 - Configuring the Server Keysize 40-32
 - Setting Up External Local CA File Storage 40-33
 - Downloading CRLs 40-35
 - Storing CRLs 40-36
 - Setting Up Enrollment Parameters 40-37
 - Adding and Enrolling Users 40-38
 - Renewing Users 40-40
 - Restoring Users 40-41
 - Removing Users 40-41
 - Revoking Certificates 40-42
 - Maintaining the Local CA Certificate Database 40-42
 - Rolling Over Local CA Certificates 40-42
 - Archiving the Local CA Server Certificate and Keypair 40-43
- Monitoring Digital Certificates 40-43
- Feature History for Certificate Management 40-45

PART 8

System Administration

CHAPTER 41

Configuring Management Access 41-1

- Configuring ASA Access for ASDM, Telnet, or SSH 41-1
 - Licensing Requirements for ASA Access for ASDM, Telnet, or SSH 41-1
 - Guidelines and Limitations 41-2
 - Configuring Telnet Access 41-3
 - Using a Telnet Client 41-3
 - Configuring SSH Access 41-4
 - Using an SSH Client 41-5
 - Configuring HTTPS Access for ASDM 41-6
- Configuring CLI Parameters 41-6
 - Licensing Requirements for CLI Parameters 41-7
 - Guidelines and Limitations 41-7
 - Configuring a Login Banner 41-7
 - Customizing a CLI Prompt 41-8

Changing the Console Timeout	41-9
Configuring ICMP Access	41-10
Information About ICMP Access	41-10
Licensing Requirements for ICMP Access	41-10
Guidelines and Limitations	41-11
Default Settings	41-11
Configuring ICMP Access	41-12
Configuring Management Access Over a VPN Tunnel	41-13
Licensing Requirements for a Management Interface	41-13
Guidelines and Limitations	41-13
Configuring a Management Interface	41-14
Configuring AAA for System Administrators	41-14
Information About AAA for System Administrators	41-14
Information About Management Authentication	41-15
Information About Command Authorization	41-16
Licensing Requirements for AAA for System Administrators	41-18
Prerequisites	41-18
Guidelines and Limitations	41-19
Default Settings	41-19
Configuring Authentication for CLI and ASDM Access	41-20
Configuring Authentication to Access Privileged EXEC Mode (the enable Command)	41-21
Configuring Authentication for the enable Command	41-22
Authenticating Users with the login Command	41-22
Limiting User CLI and ASDM Access with Management Authorization	41-23
Configuring a Password Policy for Local Database Users	41-24
Configuring the Password Policy	41-25
Changing Your Password	41-27
Configuring Command Authorization	41-27
Configuring Local Command Authorization	41-27
Viewing Local Command Privilege Levels	41-31
Configuring Commands on the TACACS+ Server	41-32
Configuring TACACS+ Command Authorization	41-33
Configuring Management Access Accounting	41-33
Viewing the Currently Logged-In User	41-34
Setting a Management Session Quota	41-35
Exchanging Keys in an SSH Session	41-35
Recovering from a Lockout	41-36
Feature History for Management Access	41-37

Managing Software and Configurations 42-1

- Upgrading the Software 42-1
 - Upgrade Path and Migrations 42-1
 - Viewing Your Current Version 42-3
 - Downloading the Software from Cisco.com 42-3
 - Upgrading a Standalone Unit 42-3
 - Upgrading a Failover Pair or ASA Cluster 42-5
 - Upgrading an Active/Standby Failover Pair 42-5
 - Upgrading an Active/Active Failover Pair 42-8
 - Upgrading an ASA Cluster 42-10
- Managing Files 42-12
 - Viewing Files in Flash Memory 42-12
 - Deleting Files from Flash Memory 42-12
 - Erasing the Flash File System 42-13
 - Configuring File Access 42-13
 - Configuring the FTP Client Mode 42-13
 - Configuring the ASA as a Secure Copy Server 42-14
 - Customizing the ASA Secure Copy Client 42-14
 - Configuring the ASA TFTP Client Path 42-16
 - Copying a File to the ASA 42-17
 - Copying a File to the Startup or Running Configuration 42-19
 - Configuring the Images and Startup Configuration to Use 42-21
 - Configuring the ASA and ASDM Images to Use 42-21
 - Configuring the File to Boot as the Startup Configuration 42-22
 - Using the ROM Monitor to Load an Image 42-22
 - Using ROM Monitor for the ASA 5500 Series 42-22
 - Using the ROM Monitor for the ASASM 42-23
 - Backing Up Configurations or Other Files 42-25
 - Backing up the Single Mode Configuration or Multiple Mode System Configuration 42-25
 - Backing Up a Context Configuration or Other File in Flash Memory 42-26
 - Backing Up a Context Configuration within a Context 42-27
 - Copying the Configuration from the Terminal Display 42-27
 - Backing Up Additional Files Using the Export and Import Commands 42-27
 - Using a Script to Back Up and Restore Files 42-28
 - Prerequisites 42-28
 - Running the Script 42-29
 - Sample Script 42-29
 - Downgrading Your Software 42-34
 - Information About Activation Key Compatibility 42-34

Performing the Downgrade	42-35
Configuring Auto Update	42-35
Information About Auto Update	42-36
Auto Update Client or Server	42-36
Auto Update Benefits	42-36
Auto Update Server Support in Failover Configurations	42-36
Guidelines and Limitations	42-39
Configuring Communication with an Auto Update Server	42-39
Configuring Client Updates as an Auto Update Server	42-41
Viewing Auto Update Status	42-42
Feature History for Software and Configurations	42-42

CHAPTER 43**Troubleshooting 43-1**

Viewing Debugging Messages	43-1
Capturing Packets	43-2
Capturing Packets in a Clustering Environment	43-4
Guidelines and Limitations	43-4
Viewing the Crash Dump	43-6
Viewing the Coredump	43-6

PART 9**Configuring Logging, SNMP, and Smart Call Home****CHAPTER 44****Configuring Logging 44-1**

Information About Logging	44-1
Logging in Multiple Context Mode	44-2
Analyzing Syslog Messages	44-2
Syslog Message Format	44-3
Severity Levels	44-3
Message Classes and Range of Syslog IDs	44-4
Filtering Syslog Messages	44-4
Using Custom Message Lists	44-5
Using Clustering	44-5
Licensing Requirements for Logging	44-5
Prerequisites for Logging	44-5
Guidelines and Limitations	44-6
Configuring Logging	44-7
Enabling Logging	44-7
Configuring an Output Destination	44-7

- Sending Syslog Messages to an External Syslog Server 44-8
- Sending Syslog Messages to the Internal Log Buffer 44-9
- Sending Syslog Messages to an E-mail Address 44-11
- Sending Syslog Messages to ASDM 44-12
- Sending Syslog Messages to the Console Port 44-12
- Sending Syslog Messages to an SNMP Server 44-12
- Sending Syslog Messages to a Telnet or SSH Session 44-13
- Creating a Custom Event List 44-14
- Generating Syslog Messages in EMBLEM Format to a Syslog Server 44-15
 - Generating Syslog Messages in EMBLEM Format to Other Output Destinations 44-15
- Changing the Amount of Internal Flash Memory Available for Logs 44-16
- Configuring the Logging Queue 44-16
- Sending All Syslog Messages in a Class to a Specified Output Destination 44-17
- Enabling Secure Logging 44-17
- Including the Device ID in Non-EMBLEM Format Syslog Messages 44-18
- Including the Date and Time in Syslog Messages 44-19
- Disabling a Syslog Message 44-19
- Changing the Severity Level of a Syslog Message 44-19
- Limiting the Rate of Syslog Message Generation 44-20
- Monitoring the Logs 44-20
- Configuration Examples for Logging 44-21
- Feature History for Logging 44-21

CHAPTER 45

Configuring SNMP 45-1

- Information About SNMP 45-1
 - Information About SNMP Terminology 45-2
 - Information About MIBs and Traps 45-3
 - SNMP Object Identifiers 45-3
 - SNMP Physical Vendor Type Values 45-5
 - Supported Tables in MIBs 45-11
 - Supported Traps (Notifications) 45-12
 - SNMP Version 3 45-15
 - SNMP Version 3 Overview 45-15
 - Security Models 45-16
 - SNMP Groups 45-16
 - SNMP Users 45-16
 - SNMP Hosts 45-16
 - Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software 45-16

Licensing Requirements for SNMP	45-17
Prerequisites for SNMP	45-17
Guidelines and Limitations	45-17
Configuring SNMP	45-18
Enabling SNMP	45-18
Configuring SNMP Traps	45-20
Configuring a CPU Usage Threshold	45-21
Configuring a Physical Interface Threshold	45-21
Using SNMP Version 1 or 2c	45-22
Using SNMP Version 3	45-23
Troubleshooting Tips	45-24
Interface Types and Examples	45-25
Monitoring SNMP	45-26
SNMP Syslog Messaging	45-27
SNMP Monitoring	45-27
Configuration Examples for SNMP	45-28
Configuration Example for SNMP Versions 1 and 2c	45-28
Configuration Example for SNMP Version 3	45-28
Where to Go Next	45-29
Additional References	45-29
RFCs for SNMP Version 3	45-29
MIBs	45-29
Application Services and Third-Party Tools	45-31
Feature History for SNMP	45-31

CHAPTER 46**Configuring NetFlow Secure Event Logging (NSEL) 46-1**

Information About NSEL	46-1
Using NSEL and Syslog Messages	46-2
Using NSEL in Clustering	46-3
Licensing Requirements for NSEL	46-4
Prerequisites for NSEL	46-4
Guidelines and Limitations	46-4
Configuring NSEL	46-5
Configuring NSEL Collectors	46-5
Configuring Flow-Export Actions Through Modular Policy Framework	46-6
Configuring Template Timeout Intervals	46-7
Changing the Time Interval for Sending Flow-Update Events to a Collector	46-8
Delaying Flow-Creation Events	46-9

- Disabling and Reenabling NetFlow-related Syslog Messages 46-9
- Clearing Runtime Counters 46-10
- Monitoring NSEL 46-10
 - NSEL Monitoring Commands 46-10
- Configuration Examples for NSEL 46-12
- Where to Go Next 46-13
- Additional References 46-13
 - Related Documents 46-14
 - RFCs 46-14
- Feature History for NSEL 46-14

CHAPTER 47

- Configuring Anonymous Reporting and Smart Call Home 47-1**
 - Information About Anonymous Reporting and Smart Call Home 47-1
 - Information About Anonymous Reporting 47-2
 - DNS Requirement 47-2
 - Anonymous Reporting and Smart Call Home Prompt 47-2
 - Information About Smart Call Home 47-3
 - Licensing Requirements for Anonymous Reporting and Smart Call Home 47-4
 - Prerequisites for Smart Call Home and Anonymous Reporting 47-4
 - Guidelines and Limitations 47-4
 - Configuring Anonymous Reporting and Smart Call Home 47-5
 - Configuring Anonymous Reporting 47-6
 - Configuring Smart Call Home 47-6
 - Enabling Smart Call Home 47-6
 - Declaring and Authenticating a CA Trust Point 47-7
 - Subscribing to Alert Groups 47-8
 - Optional Configuration Procedures 47-15
 - Monitoring Anonymous Reporting and Smart Call Home 47-22
 - Configuration Example for Smart Call Home 47-23
 - Feature History for Anonymous Reporting and Smart Call Home 47-24

PART 10

Reference

APPENDIX 48

- Using the Command-Line Interface 48-1**
 - Firewall Mode and Security Context Mode 48-1
 - Command Modes and Prompts 48-2
 - Syntax Formatting 48-3

Abbreviating Commands	48-3
Command-Line Editing	48-3
Command Completion	48-4
Command Help	48-4
Viewing the Running Configuration	48-4
Filtering show and more Command Output	48-5
Command Output Paging	48-5
Adding Comments	48-6
Text Configuration Files	48-6
How Commands Correspond with Lines in the Text File	48-6
Command-Specific Configuration Mode Commands	48-6
Automatic Text Entries	48-7
Line Order	48-7
Commands Not Included in the Text Configuration	48-7
Passwords	48-7
Multiple Security Context Files	48-7
Supported Character Sets	48-8

APPENDIX 49

Addresses, Protocols, and Ports	49-1
IPv4 Addresses and Subnet Masks	49-1
Classes	49-1
Private Networks	49-2
Subnet Masks	49-2
Determining the Subnet Mask	49-3
Determining the Address to Use with the Subnet Mask	49-3
IPv6 Addresses	49-5
IPv6 Address Format	49-5
IPv6 Address Types	49-6
Unicast Addresses	49-6
Multicast Address	49-8
Anycast Address	49-9
Required Addresses	49-10
IPv6 Address Prefixes	49-10
Protocols and Applications	49-11
TCP and UDP Ports	49-11
Local Ports and Protocols	49-14
ICMP Types	49-15



About This Guide

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

- [Document Objectives, page i](#)
- [Related Documentation, page i](#)
- [Conventions, page i](#)
- [Obtaining Documentation and Submitting a Service Request, page ii](#)

Document Objectives

The purpose of this guide is to help you configure general operations for the ASA using the command-line interface. This guide does not cover every feature, but describes only the most common configuration scenarios.

You can also configure and monitor the ASA by using ASDM, a web-based GUI application. ASDM includes configuration wizards to guide you through some common configuration scenarios, and online help for less common scenarios.

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

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 .

<i>italic font</i>	Document titles, new or emphasized terms, and arguments for which you supply values are in <i>italic font</i> .
[]	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.
<code>courier font</code>	Terminal sessions and information the system displays appear in <code>courier font</code> .
<code>courier bold font</code>	Commands and keywords and user-entered text appear in <code>courier bold font</code> .
<i><code>courier italic font</code></i>	Arguments for which you supply values are in <i><code>courier italic font</code></i> .
< >	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.

This chapter includes the following sections:

- [Hardware and Software Compatibility, page 1-1](#)
- [VPN Compatibility, page 1-1](#)
- [New Features, page 1-2](#)
- [Firewall Functional Overview, page 1-14](#)
- [VPN Functional Overview, page 1-18](#)
- [Security Context Overview, page 1-19](#)
- [ASA Clustering Overview, page 1-19](#)

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/asamatrix.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\(5\)](#), page 1-2
- [New Features in ASA 9.1\(4\)](#), page 1-3
- [New Features in ASA 9.1\(3\)](#), page 1-5
- [New Features in ASA 9.1\(2\)](#), page 1-7
- [New Features in ASA 9.1\(1\)](#), page 1-13


Note

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

New Features in ASA 9.1(5)

Released: March 31, 2014

[Table 1-1](#) lists the new features for ASA Version 9.1(5).

Table 1-1 **New Features for ASA Version 9.1(5)**

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 introduced the following commands: ssh pubkey-chain , server (ssh pubkey-chain) , key-string , key-hash , ssh stricthostkeycheck . We modified the following command: copy scp .
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 command: aaa authorization exec .
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 commands: asp rule-engine transactional-commit , show running-config asp rule-engine transactional-commit , clear configure asp rule-engine transactional-commit .
Monitoring Features	

Table 1-1 *New Features for ASA Version 9.1(5) (continued)*

Feature	Description
SNMP hosts, host groups, and user lists	<p>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.</p> <p>We introduced or modified the following commands: snmp-server host-group, snmp-server user-list, show running-config snmp-server, clear configure snmp-server.</p>
Remote Access Features	
AnyConnect DTLS Single Session Performance Improvement	<p>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.</p> <p>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.</p> <p>We did not modify any commands.</p>
Webtype ACL enhancements	<p>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.</p> <p>We did not modify any commands.</p>

New Features in ASA 9.1(4)

Released: December 9, 2013

[Table 1-2](#) lists the new features for ASA Version 9.1(4).

Table 1-2 *New Features for ASA Version 9.1(4)*

Feature	Description
Remote Access Features	
HTML5 WebSocket proxying	<p>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.</p> <p>We did not modify any commands.</p>

Table 1-2 New Features for ASA Version 9.1(4) (continued)

Feature	Description
Inner IPv6 for IKEv2	<p>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.</p> <p>Note This feature requires AnyConnect Client Version 3.1.05 or later.</p> <p>Output of the show ipsec sa and show vpn-sessiondb detail anyconnect commands has been updated to reflect the assigned IPv6 address, and to indicate the GRE Transport Mode security association when doing IKEv2 dual traffic.</p> <p>The vpn-filter command must now be used for both IPv4 and IPv6 ACLs. If the deprecated ipv6-vpn-filter command is used to configure IPv6 ACLs the connection will be terminated.</p>
Mobile Devices running Citrix Server Mobile have additional connection options	<p>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.</p> <p>We introduced the application-type command to configure the default tunnel group for VDI connections when a Citrix Receiver user does not choose a tunnel-group. A none action was added to the vdj command to disable VDI configuration for a particular group policy or user.</p>
Split-tunneling supports exclude ACLs	<p>Split-tunneling of VPN traffic has been enhanced to support both exclude and include ACLs. Exclude ACLs were previously ignored.</p> <p>Note This feature requires AnyConnect Client Version 3.1.03103 or later.</p> <p>We did not modify any commands.</p>
High Availability and Scalability Features	
ASA 5500-X support for clustering	<p>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.</p> <p>We did not modify any commands.</p>
Improved VSS and vPC support for health check monitoring	<p>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.</p> <p>We modified the following command: health-check [vss-enabled]</p>

Table 1-2 ***New Features for ASA Version 9.1(4) (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 commands.
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 command: health-check [vss-enabled]
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 commands: show ip address dhcp lease proxy , show ip address dhcp lease summary , and show ip address dhcp lease server .
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: <ul style="list-style-type: none"> • Creating an AK47 instance. • Destroying an AK47 instance. • Generating a 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)

Released: September 18, 2013

Table 1-3 lists the new features for ASA Version 9.1(3).

Table 1-3 New Features for ASA Version 9.1(3)

Feature	Description
Module Features	
Support for the ASA CX module in multiple context mode	<p>You can now configure ASA CX service policies per context on the ASA.</p> <p>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.</p> <p>Requires ASA CX 9.2(1) or later.</p> <p>We did not modify any commands.</p>
ASA 5585-X with SSP-40 and -60 support for the ASA CX SSP-40 and -60	<p>ASA CX SSP-40 and -60 modules can be used with the matching level ASA 5585-X with SSP-40 and -60.</p> <p>Requires ASA CX 9.2(1) or later.</p> <p>We did not modify any commands.</p>
Filtering packets captured on the ASA CX backplane	<p>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.</p> <p>Requires ASA CX 9.2(1) or later.</p> <p>We modified the following command: capture interface asa_dataplane.</p>
Monitoring Features	
Ability to view top 10 memory users	<p>You can now view the top bin sizes allocated and the top 10 PCs for each allocated bin size. Previously, you had to enter multiple commands to see this information (the show memory detail command and the show memory binsize command); the new command provides for quicker analysis of memory issues.</p> <p>We introduced the following command: show memory top-usage.</p> <p><i>Also available in 8.4(6).</i></p>

Table 1-3 New Features for ASA Version 9.1(3) (continued)

Feature	Description
Smart Call Home	<p>We added a new type of Smart Call Home message to support ASA clustering.</p> <p>A Smart Call Home clustering message is sent for only the following three events:</p> <ul style="list-style-type: none"> • When a unit joins the cluster • When a unit leaves the cluster • When a cluster unit becomes the cluster master <p>Each message that is sent includes the following information:</p> <ul style="list-style-type: none"> • The active cluster member count • The output of the show cluster info command and the show cluster history command on the cluster master <p>We modified the following commands: show call-home, show running-config call-home.</p> <p><i>Also available in 9.0(3).</i></p>
Remote Access Features	
user-storage value command password is now encrypted in show commands	<p>The password in the user-storage value command is now encrypted when you enter show running-config.</p> <p>We modified the following command: user-storage value.</p> <p><i>Also available in 8.4(6).</i></p>

New Features in ASA 9.1(2)

Released: May 14, 2013

Table 1-4 lists the new features for ASA Version 9.1(2).



Note

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

Table 1-4 New Features for ASA Version 9.1(2)

Feature	Description
Certification Features	
FIPS and Common Criteria certifications	<p>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.</p> <p>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.</p>
Encryption Features	

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	<p>Instead of using the proprietary encryption for the failover key (the failover key command), you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption.</p> <p>Note Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.</p> <p>We introduced or modified the following commands: failover ipsec pre-shared-key, show vpn-sessiondb.</p>
Additional ephemeral Diffie-Hellman ciphers for SSL encryption	<p>The ASA now supports the following ephemeral Diffie-Hellman (DHE) SSL cipher suites:</p> <ul style="list-style-type: none"> • DHE-AES128-SHA1 • DHE-AES256-SHA1 <p>These cipher suites are specified in RFC 3268, <i>Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS)</i>.</p> <p>When supported by the client, DHE is the preferred cipher because it provides Perfect Forward Secrecy. See the following limitations:</p> <ul style="list-style-type: none"> • DHE is not supported on SSL 3.0 connections, so make sure to also enable TLS 1.0 for the SSL server. <pre>!! set server version hostname(config)# ssl server-version tlsv1 sslv3 !! set client version hostname(config) # ssl client-version any</pre> <ul style="list-style-type: none"> • 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. <p>We modified the following command: ssl encryption.</p> <p><i>Also available in 8.4(4.1).</i></p>
Management Features	
Support for administrator password policy when using the local database	<p>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.</p> <p>We introduced the following commands: change-password, password-policy lifetime, password-policy minimum changes, password-policy minimum-length, password-policy minimum-lowercase, password-policy minimum-uppercase, password-policy minimum-numeric, password-policy minimum-special, password-policy authenticate enable, clear configure password-policy, show running-config password-policy.</p> <p><i>Also available in 8.4(4.1).</i></p>

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
Support for SSH public key authentication	<p>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).</p> <p>We introduced the following commands: ssh authentication.</p> <p><i>Also available in 8.4(4.1); PKF key format support is only in 9.1(2).</i></p>
AES-CTR encryption for SSH	The SSH server implementation in the ASA now supports AES-CTR mode encryption.
Improved SSH rekey interval	<p>An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic.</p> <p>We introduced the following command: show ssh sessions detail.</p>
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	<p>Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported.</p> <p>We introduced the following command: ssh key-exchange.</p> <p><i>Also available in 8.4(4.1).</i></p>
Support for a maximum number of management sessions	<p>You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions.</p> <p>We introduced the following commands: quota management-session, show running-config quota management-session, show quota management-session.</p> <p><i>Also available in 8.4(4.1).</i></p>
The default Telnet password was removed	<p>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 (the aaa authentication telnet console command).</p> <p>Formerly, when you cleared the password, the ASA restored the default of “cisco.” Now when you clear the password, the password is removed.</p> <p>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.</p> <p>We modified the following command: passwd.</p> <p><i>Also available in 9.0(2).</i></p>
Platform Features	
Support for Power-On Self-Test (POST)	<p>The ASA runs its power-on self-test at boot time even if it is not running in FIPS 140-2-compliant mode.</p> <p>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).</p>
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	<p>Support for SHA-512 image integrity checking was added.</p> <p>We modified the following command: verify.</p> <p><i>Also available in 8.4(4.1).</i></p>

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
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.
CPU profile enhancements	<p>The cpu profile activate command now supports the following:</p> <ul style="list-style-type: none"> • Delayed start of the profiler until triggered (global or specific thread CPU%) • Sampling of a single thread <p>We modified the following command: cpu profile activate [<i>n-samples</i>] [sample-process <i>process-name</i>] [trigger cpu-usage <i>cpu%</i>] [<i>process-name</i>].</p> <p><i>Also available in 8.4(6).</i></p>
DHCP Features	
DHCP relay servers per interface (IPv4 only)	<p>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.</p> <p>We introduced or modified the following commands: dhcprelay server (interface config mode), clear configure dhcprelay, show running-config dhcprelay.</p>
DHCP trusted interfaces	<p>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.</p> <p>We introduced or modified the following commands: dhcprelay information trusted, dhcprelay informarion trust-all, show running-config dhcprelay.</p>
Module Features	
ASA 5585-X support for network modules	<p>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:</p> <ul style="list-style-type: none"> • ASA 4-port 10G Network Module • ASA 8-port 10G Network Module • ASA 20-port 1G Network Module <p><i>Also available in 8.4(4.1).</i></p>
ASA 5585-X DC power supply support	<p>Support was added for the ASA 5585-X DC power supply.</p> <p><i>Also available in 8.4(5).</i></p>

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
Support for ASA CX monitor-only mode for demonstration purposes	<p>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.</p> <p>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.</p> <p>We modified or introduced the following commands: cxsc {fail-close fail-open} monitor-only, traffic-forward cxsc monitor-only.</p>
Support for the ASA CX module and NAT 64	<p>You can now use NAT 64 in conjunction with the ASA CX module.</p> <p>We did not modify any commands.</p>
NetFlow Features	
Support for NetFlow flow-update events and an expanded set of NetFlow templates	<p>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.</p> <p>Two new fields were added for IPv6 translation support.</p> <p>Several NetFlow field IDs were changed to their IPFIX equivalents.</p> <p>For more information, see the <i>Cisco ASA Implementation Note for NetFlow Collectors</i>.</p>
Firewall Features	
EtherType ACL support for IS-IS traffic (transparent firewall mode)	<p>In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL.</p> <p>We modified the following command: access-list ethertype {permit deny} is-is.</p> <p><i>Also available in 8.4(5).</i></p>
Decreased the half-closed timeout minimum value to 30 seconds	<p>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.</p> <p>We modified the following commands: set connection timeout half-closed, timeout half-closed.</p>
Remote Access Features	

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
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 command: crypto ikev1 limit .
	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 introduced the following command: crypto ipsec ikev2 sa-strength-enforcement .
	For Site-to-Site, IPsec data-based rekeying can be disabled. We modified the following command: crypto ipsec security-association .
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. <i>Also available in 8.4(5).</i>
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: <ul style="list-style-type: none"> • Internet Explorer 10 (desktop only) • Firefox (all supported Windows 8 versions) • Chrome (all supported Windows 8 versions) See the following limitations: <ul style="list-style-type: none"> • Internet Explorer 10: <ul style="list-style-type: none"> – 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. <i>Also available in 9.0(2).</i>

Table 1-4 New Features for ASA Version 9.1(2) (continued)

Feature	Description
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: <ul style="list-style-type: none"> Secure Desktop (Vault) is not supported with Windows 8. <i>Also available in 9.0(2).</i>
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. <i>Also available in 8.4(5).</i>
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 introduced or modified the following commands: flow-export active refresh-interval , flow-export event-type . <i>Also available in 8.4(5).</i>

New Features in ASA 9.1(1)

Released: December 3, 2012

Table 1-5 lists the new features for ASA Version 9.1(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 1-5 New Features for ASA Version 9.1(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 modified the following commands: session cxsc , show module cxsc , sw-module cxsc .

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 1-14](#)
- [Firewall Mode Overview, page 1-17](#)
- [Stateful Inspection Overview, page 1-17](#)

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 Lists, page 1-15](#)
- [Applying NAT, page 1-15](#)
- [Protecting from IP Fragments, page 1-15](#)
- [Using AAA for Through Traffic, page 1-15](#)
- [Applying HTTP, HTTPS, or FTP Filtering, page 1-15](#)
- [Applying Application Inspection, page 1-15](#)
- [Sending Traffic to a Module, page 1-15](#)
- [Applying QoS Policies, page 1-16](#)
- [Applying Connection Limits and TCP Normalization, page 1-16](#)
- [Enabling Threat Detection, page 1-16](#)
- [Enabling the Botnet Traffic Filter, page 1-16](#)
- [Configuring Cisco Unified Communications, page 1-16](#)

Permitting or Denying Traffic with Access Lists

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

Applying NAT

Some of the benefits of NAT include the following:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the local addresses from other networks, so attackers cannot learn the real address of a host.
- NAT can resolve IP routing problems by supporting overlapping IP addresses.

Protecting from IP Fragments

The 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 use 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 2-5](#)
- [Verifying the Module Installation, page 2-6](#)
- [Assigning VLANs to the ASA Services Module, page 2-7](#)
- [Using the MSFC as a Directly Connected Router \(SVIs\), page 2-10](#)
- [Configuring the Switch for ASA Failover, page 2-11](#)
- [Resetting the ASA Services Module, page 2-12](#)
- [Monitoring the ASA Services Module, page 2-12](#)
- [Feature History for the Switch for Use with the ASA Services Module, page 2-15](#)

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.

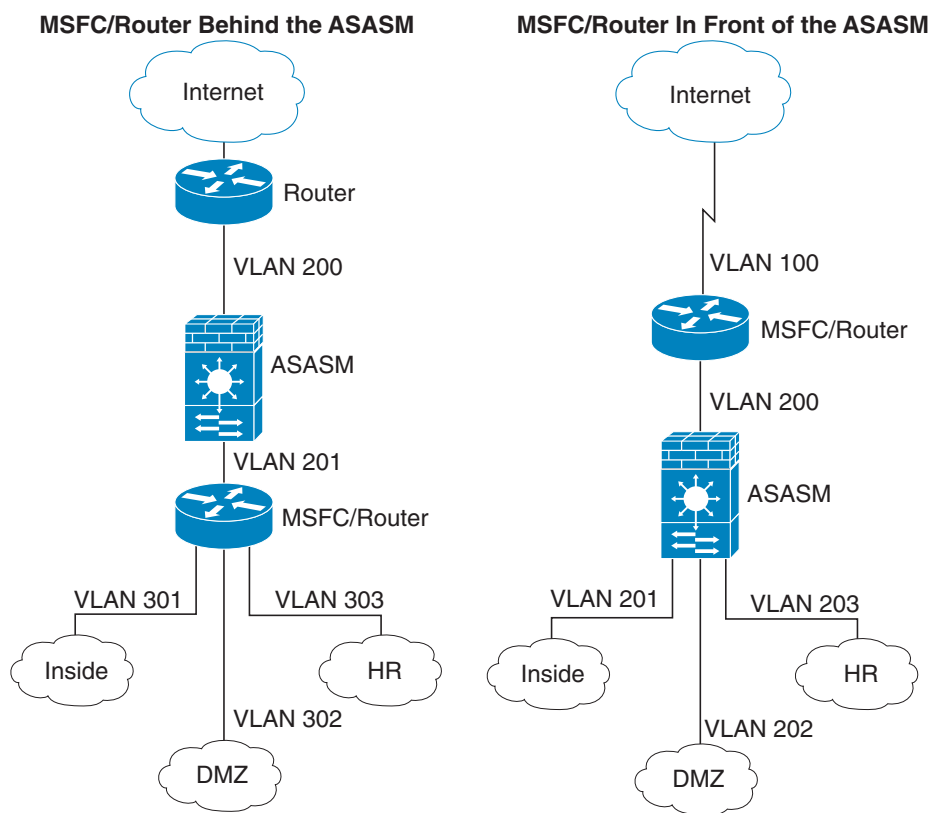
The switch includes a switching processor (the supervisor) and a router (the MSFC). Although you need the MSFC as part of your system, you do not have to use it. If you choose to do so, you can assign one or more VLAN interfaces to the MSFC. You can alternatively use external routers instead of the MSFC.

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/asamatrix.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.

You might need to bypass the ASASM in some network scenarios. [Figure 2-3](#) 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 2-3 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. 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.

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 2-7](#).
- 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 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
<code>show module [switch {1 2}] [mod-num all]</code>	Displays module information. For a switch in a VSS, enter the <code>switch</code> keyword.
Example: <code>Router# show module 1</code>	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 2-6.

Detailed Steps

	Command	Purpose
Step 1	<pre>firewall vlan-group firewall_group vlan_range</pre> <p>Example: Router(config)# firewall vlan-group 1 55-57</p>	<p>Assigns VLANs to a firewall group.</p> <p>The <i>firewall_group</i> argument is an integer. The <i>vlan_range</i> argument can be one or more VLANs (2 to 1001) identified in one of the following ways:</p> <ul style="list-style-type: none"> • A single number (<i>n</i>) • A range (<i>n-x</i>) <p>Separate numbers or ranges by commas, as shown in the following example:</p> <p>5,7-10,13,45-100</p>
Step 2	<pre>firewall [switch {1 2}] module slot vlan-group firewall_group</pre> <p>Example: Router(config)# firewall module 5 vlan-group 1</p>	<p>Assigns the firewall groups to the ASASM.</p> <p>For a switch in a VSS, enter the switch argument.</p> <p>To view the slots where the ASASM is installed, enter the show module command.</p> <p>The <i>firewall_group</i> argument is one or more group numbers, which can be one of the following:</p> <ul style="list-style-type: none"> • A single number (<i>n</i>) • A range (<i>n-x</i>) <p>Separate numbers or ranges by commas, as shown in the following example:</p> <p>5,7-10</p>

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 <i>vlan_number</i> Example: Router(config)# interface vlan 55	Adds a VLAN interface to the MSFC.
Step 3	ip address <i>address mask</i> 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 2-11](#)
- [Adding a Trunk Between a Primary Switch and Secondary Switch, page 2-11](#)
- [Ensuring Compatibility with Transparent Firewall Mode, page 2-11](#)
- [Enabling Autostate Messaging for Rapid Link Failure Detection, page 2-11](#)

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 2-11](#).

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
<code>firewall autostate</code>	Enables autostate messaging in Cisco IOS software. Autostate messaging is disabled by default.
Example: Router(config)# <code>firewall autostate</code>	

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
<code>hw-module [switch {1 2}] module slot reset</code>	Resets the ASASM.
Example: Router# <code>hw-module module 9 reset</code>	For a switch in a VSS, enter the switch argument. The <i>slot</i> 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
<code>show firewall module [mod-num] state</code>	Verifies the state of the ASA.
<code>show firewall module [mod-num] traffic</code>	Verifies that traffic is flowing through the ASA.

Command	Purpose
<code>show firewall module [mod-num] version</code>	Shows the software version of the ASA.
<code>show firewall multiple-vlan-interfaces</code>	Indicates the status of multiple VLAN interfaces (enabled or disabled).
<code>show firewall vlan-group</code>	Displays all configured VLAN groups.
<code>show interface vlan</code>	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 3-6](#)
- [Configuring ASDM Access for the ASA Services Module, page 3-12](#)
- [Starting ASDM, page 3-14](#)
- [Factory Default Configurations, page 3-18](#)
- [Working with the Configuration, page 3-24](#)
- [Applying Configuration Changes to Connections, page 3-28](#)
- [Reloading the ASA, page 3-29](#)

Accessing the Appliance Command-Line Interface

For initial configuration, access the CLI directly from the console port. Later, you can configure remote access using Telnet or SSH according to [Chapter 41, “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 6, “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 13-1 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. 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 6, “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 41-1.

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:	
	<p>(Available for initial access.)</p> <pre>service-module session [switch {1 2}] slot <i>number</i></pre> <p>Example: Router# service-module session slot 3 ciscoasa></p>	<p>From the switch CLI, enter this command to gain console access to the ASASM.</p> <p>For a switch in a VSS, enter the switch argument.</p> <p>To view the module slot numbers, enter the show module command at the switch prompt.</p> <p>You access user EXEC mode.</p>
	<p>(Available after you configure a login password.)</p> <pre>session [switch {1 2}] slot <i>number</i> processor 1</pre> <p>You are prompted for the login password: ciscoasa passwd:</p> <p>Example: Router# session slot 3 processor 1 ciscoasa passwd: cisco ciscoasa></p>	<p>From the switch CLI, enter this command to Telnet to the ASASM over the backplane.</p> <p>For a switch in a VSS, enter the switch argument.</p> <p>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.</p> <p>To view the module slot numbers, enter the show module command at the switch prompt.</p> <p>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.</p> <p>You access user EXEC mode.</p>
Step 2	<pre>enable</pre> <p>Example: ciscoasa> enable Password: ciscoasa#</p>	<p>Accesses privileged EXEC mode, which is the highest privilege level.</p> <p>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 13-1.</p> <p>To exit privileged EXEC mode, enter the disable, exit, or quit command.</p>
Step 3	<pre>configure terminal</pre> <p>Example: ciscoasa# configure terminal ciscoasa(config)#</p>	<p>Accesses global configuration mode.</p> <p>To exit global configuration mode, enter the disable, exit, or quit command.</p>

Logging Out of a Console Session

This section includes the following topics:

- [Logging Out, page 3-5](#)
- [Killing an Active Console Connection, page 3-5](#)

Logging Out

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

To kill another user's active connection, which may have been unintentionally left open, see the [“Killing an Active Console Connection” section on page 3-5](#).

Detailed Steps

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

Ctrl-Shift-6, x

You return to the switch prompt:

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



Note

Shift-6 on US and UK keyboards issues the caret (^) character. If you have a different keyboard and cannot issue the caret (^) character as a standalone character, you can temporarily or permanently change the escape character to a different character. 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:

- [Accessing ASDM Using the Factory Default Configuration, page 3-6](#)
- [Customizing ASDM Access \(ASA 5505\), page 3-7](#)
- [Customizing ASDM Access \(ASA 5510 and Higher\), page 3-10](#)

Accessing ASDM Using the Factory Default Configuration

With a factory default configuration (see the “[Factory Default Configurations](#)” section on page 3-18), 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 6-16. 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 the management IP address
- You want to change to transparent firewall mode

See also the sample configurations in the [“ASA 5505 Default Configuration”](#) section on page 3-20.



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) <code>firewall transparent</code> Example: <code>ciscoasa(config)# firewall transparent</code>	Enables transparent firewall mode. This command clears your configuration. See the “Setting the Firewall Mode” section on page 5-9 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode:	

Command	Purpose
<p>Routed mode:</p> <pre>interface vlan <i>number</i> nameif <i>name</i> security-level <i>level</i> ip address <i>ip_address</i> [<i>mask</i>]</pre> <p>Example:</p> <pre>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</pre>	<p>Configures an interface in routed mode. The security-level is a number between 1 and 100, where 100 is the most secure.</p>
<p>Transparent mode:</p> <pre>interface bvi <i>number</i> ip address <i>ip_address</i> [<i>mask</i>]</pre> <pre>interface vlan <i>number</i> bridge-group <i>number</i> nameif <i>name</i> security-level <i>level</i></pre> <p>Example:</p> <pre>ciscoasa(config)# interface bvi 1 ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0</pre> <pre>ciscoasa(config)# interface vlan 1 ciscoasa(config-if)# bridge-group 1 ciscoasa(config-if)# nameif inside ciscoasa(config-if)# security-level 100</pre>	<p>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.</p>
<p>Step 3</p> <pre>interface ethernet 0/1 switchport access vlan <i>number</i> no shutdown</pre> <p>Example:</p> <pre>ciscoasa(config)# interface ethernet 0/1 ciscoasa(config-if)# switchport access vlan 1 ciscoasa(config-if)# no shutdown</pre>	<p>Enables the management switchport and assigns it to the management VLAN.</p>
<p>Step 4</p> <pre>dhcpd address <i>ip_address-ip_address</i> <i>interface_name</i> dhcpd enable <i>interface_name</i></pre> <p>Example:</p> <pre>ciscoasa(config)# dhcpd address 192.168.1.5-192.168.1.254 inside ciscoasa(config)# dhcpd enable inside</pre>	<p>Sets the DHCP pool for the management network. Make sure you do not include the VLAN interface address in the range.</p> <p>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.</p>
<p>Step 5</p> <pre>http server enable</pre> <p>Example:</p> <pre>ciscoasa(config)# http server enable</pre>	<p>Enables the HTTP server for ASDM.</p>

	Command	Purpose
Step 6	<code>http ip_address mask interface_name</code> Example: ciscoasa(config)# http 192.168.1.0 255.255.255.0 inside	Allows the management host(s) to access ASDM.
Step 7	<code>write memory</code> 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 the management IP address
- 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” section on page 5-9 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.

	Command	Purpose
Step 6	<code>http ip_address mask interface_name</code> Example: ciscoasa(config)# http 192.168.1.0 255.255.255.0 management	Allows the management host(s) to access ASDM.
Step 7	<code>write memory</code> Example: ciscoasa(config)# write memory	Saves the configuration.
Step 8	(Optional) <code>mode multiple</code> 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 6, “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 2-7.](#)
- 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 ” section on page 5-9 for more information.
Step 2	Do one of the following to configure a management interface, depending on your mode: Routed mode: interface vlan <i>number</i> ip address <i>ip_address</i> [<i>mask</i>] nameif <i>name</i> security-level <i>level</i> 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 <i>number</i> ip address <i>ip_address</i> [<i>mask</i>] interface vlan <i>number</i> bridge-group <i>bvi_number</i> nameif <i>name</i> security-level <i>level</i> 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 <i>ip_address-ip_address</i> <i>interface_name</i> dhcpd enable <i>interface_name</i> 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.

	Command	Purpose
Step 4	(For remote management hosts) <code>route management_ifc management_host_ip mask gateway_ip 1</code> Example: <code>ciscoasa(config)# route management 10.1.1.0 255.255.255.0 192.168.1.50</code>	Configures a route to the management hosts.
Step 5	<code>http server enable</code> Example: <code>ciscoasa(config)# http server enable</code>	Enables the HTTP server for ASDM.
Step 6	<code>http ip_address mask interface_name</code> Example: <code>ciscoasa(config)# http 192.168.1.0 255.255.255.0 management</code>	Allows the management host to access ASDM.
Step 7	<code>write memory</code> Example: <code>ciscoasa(config)# write memory</code>	Saves the configuration.
Step 8	(Optional) <code>mode multiple</code> Example: <code>ciscoasa(config)# mode multiple</code>	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 6, “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 3-6 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.
-

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 6, “Configuring Multiple Context Mode,”](#) for more information about multiple context mode. See [Chapter 5, “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 3-20](#)
- [ASA 5510 and Higher Default Configuration, page 3-24](#)

Restoring the Factory Default Configuration

This section describes how to restore the factory default configuration.

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

	Command	Purpose
Step 1	<pre>configure factory-default [ip_address [mask]]</pre> <p>Example:</p> <pre>ciscoasa(config)# configure factory-default 10.1.1.1 255.255.255.0</pre>	<p>Restores the factory default configuration.</p> <p>If you specify the <i>ip_address</i>, 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 address command range consists of addresses within the subnet that you specify.</p> <p>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.</p>
Step 2	<pre>write memory</pre> <p>Example:</p> <pre>active(config)# write memory</pre>	<p>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.</p>

What to Do Next

See the “[Working with the Configuration](#)” section on page 3-24 to start configuring the ASA.

ASA 5505 Default Configuration

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

- [ASA 5505 Routed Mode Default Configuration, page 3-20](#)
- [ASA 5505 Transparent Mode Sample Configuration, page 3-22](#)


ASA 5505 Routed Mode Default Configuration

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

- Interfaces—Inside (VLAN 1) and outside (VLAN 2).
- Switchports enabled and assigned—Ethernet 0/1 through 0/7 switch ports assigned to inside. Ethernet 0/0 assigned to outside.
- IP addresses— Outside address from DHCP; inside address set manually to 192.168.1.1/24.
- Network address translation (NAT)—All inside IP addresses are translated when accessing the outside using interface PAT.

- Traffic flow—IPv4 and IPv6 traffic allowed from inside to outside (this behavior is implicit on the ASA). Outside users are prevented from accessing the inside.
- DHCP server—Enabled for inside hosts, so a PC connecting to the inside interface receives an address between 192.168.1.5 and 192.168.1.254. DNS, WINS, and domain information obtained from the DHCP client on the outside interface is passed to the DHCP clients on the inside interface.
- Default route—Derived from DHCP.
- ASDM access—Inside hosts allowed.

Figure 3-1 ASA 5505 Routed Mode



The configuration consists of the following commands:

```
interface Ethernet 0/0
  switchport access vlan 2
  no shutdown
interface Ethernet 0/1
  switchport access vlan 1
  no shutdown
interface Ethernet 0/2
  switchport access vlan 1
  no shutdown
interface Ethernet 0/3
  switchport access vlan 1
  no shutdown
interface Ethernet 0/4
  switchport access vlan 1
  no shutdown
interface Ethernet 0/5
  switchport access vlan 1
  no shutdown
interface Ethernet 0/6
  switchport access vlan 1
  no shutdown
interface Ethernet 0/7
  switchport access vlan 1
  no shutdown
interface vlan2
  nameif outside
  no shutdown
```

```
ip address dhcp setroute
interface vlan1
  nameif inside
  ip address 192.168.1.1 255.255.255.0
  security-level 100
  no shutdown
object network obj_any
  subnet 0 0
  nat (inside,outside) dynamic interface
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.5-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
```

**Note**

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

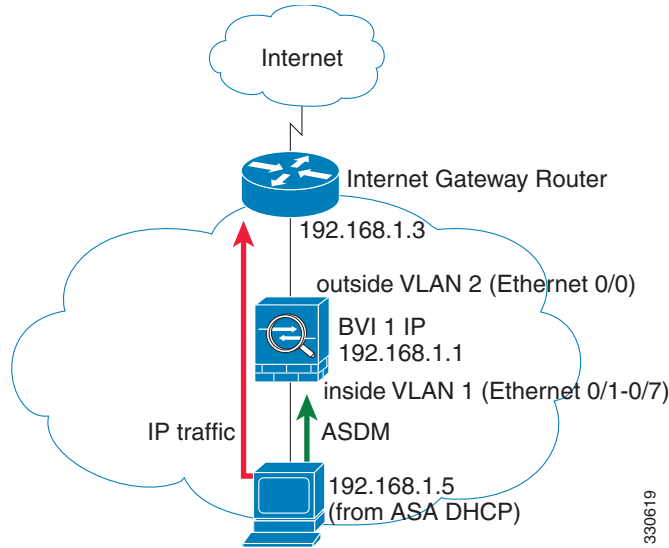
```
policy-map global_policy
  class inspection_default
    inspect icmp
```

ASA 5505 Transparent Mode Sample Configuration

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

- IP addresses—The IP addresses configured should be changed to match the network to which you are connecting.
- Static routes—For some kinds of traffic, static routes are required. See the [“MAC Address vs. Route Lookups”](#) section on page 5-6.

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
```

Working with the Configuration

This section describes how to work with the configuration. The ASA loads the configuration from a text file, called the startup configuration. This file resides by default as a hidden file in internal flash memory. You can, however, specify a different path for the startup configuration. (For more information, see [Chapter 42, “Managing Software and Configurations.”](#))

When you enter a command, the change is made only to the running configuration in memory. You must manually save the running configuration to the startup configuration for your changes to remain after a reboot.

The information in this section applies to both single and multiple security contexts, except where noted. Additional information about contexts is in [Chapter 6, “Configuring Multiple Context Mode.”](#)

This section includes the following topics:

- [Saving Configuration Changes, page 3-25](#)
- [Copying the Startup Configuration to the Running Configuration, page 3-26](#)

- [Viewing the Configuration, page 3-27](#)
- [Clearing and Removing Configuration Settings, page 3-27](#)
- [Creating Text Configuration Files Offline, page 3-28](#)

Saving Configuration Changes

This section describes how to save your configuration and includes the following topics:

- [Saving Configuration Changes in Single Context Mode, page 3-25](#)
- [Saving Configuration Changes in Multiple Context Mode, page 3-25](#)

Saving Configuration Changes in Single Context Mode

To save the running configuration to the startup configuration, enter the following command:

Command	Purpose
<code>write memory</code>	Saves the running configuration to the startup configuration.
Example: <code>ciscoasa# write memory</code>	Note The <code>copy running-config startup-config</code> command is equivalent to the <code>write memory</code> command.

Saving Configuration Changes in Multiple Context Mode

You can save each context (and system) configuration separately, or you can save all context configurations at the same time. This section includes the following topics:

- [Saving Each Context and System Separately, page 3-25](#)
- [Saving All Context Configurations at the Same Time, page 3-26](#)

Saving Each Context and System Separately

To save the system or context configuration, enter the following command within the system or context:

Command	Purpose
<code>write memory</code>	Saves the running configuration to the startup configuration.
Example: <code>ciscoasa# write memory</code>	For multiple context mode, context startup configurations can reside on external servers. In this case, the ASA saves the configuration back to the server you identified in the context URL, except for an HTTP or HTTPS URL, which do not let you save the configuration to the server. Note The <code>copy running-config startup-config</code> command is equivalent to the <code>write memory</code> command.

Saving All Context Configurations at the Same Time

To save all context configurations at the same time, as well as the system configuration, enter the following command in the system execution space:

Command	Purpose
<pre>write memory all [/noconfirm]</pre> <p>Example: <pre>ciscoasa# write memory all /noconfirm</pre></p>	<p>Saves the running configuration to the startup configuration for all contexts and the system configuration.</p> <p>If you do not enter the /noconfirm keyword, you see the following prompt: <pre>Are you sure [Y/N]:</pre></p> <p>After you enter Y, the ASA saves the system configuration and each context. Context startup configurations can reside on external servers. In this case, the ASA saves the configuration back to the server you identified in the context URL, except for an HTTP or HTTPS URL, which do not let you save the configuration to the server.</p>

After the ASA saves each context, the following message appears:

```
'Saving context 'b' ... ( 1/3 contexts saved ) '
```

Sometimes, a context is not saved because of an error. See the following information for errors:

- For contexts that are not saved because of low memory, the following message appears:

```
The context 'context a' could not be saved due to Unavailability of resources
```
- For contexts that are not saved because the remote destination is unreachable, the following message appears:

```
The context 'context a' could not be saved due to non-reachability of destination
```

- For contexts that are not saved because the context is locked, the following message appears:

```
Unable to save the configuration for the following contexts as these contexts are locked.
context 'a' , context 'x' , context 'z' .
```

A context is only locked if another user is already saving the configuration or in the process of deleting the context.

- For contexts that are not saved because the startup configuration is read-only (for example, on an HTTP server), the following message report is printed at the end of all other messages:

```
Unable to save the configuration for the following contexts as these contexts have read-only config-urls:
context 'a' , context 'b' , context 'c' .
```

- For contexts that are not saved because of bad sectors in the flash memory, the following message appears:

```
The context 'context a' could not be saved due to Unknown errors
```

Copying the Startup Configuration to the Running Configuration

Copy a new startup configuration to the running configuration using one of the following options.

Command	Purpose
<code>copy startup-config running-config</code>	Merges the startup 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.
<code>reload</code>	Reloads the ASA, which loads the startup configuration and discards the running configuration.
<code>clear configure all</code> <code>copy startup-config running-config</code>	Loads the startup configuration and discards the running configuration without requiring a reload.

Viewing the Configuration

The following commands let you view the running and startup configurations.

Command	Purpose
<code>show running-config</code>	Views the running configuration.
<code>show running-config <i>command</i></code>	Views the running configuration of a specific command.
<code>show startup-config</code>	Views the startup configuration.

Clearing and Removing Configuration Settings

To erase settings, enter one of the following commands.

Command	Purpose
<code>clear configure <i>configurationcommand</i></code> <code>[<i>level2configurationcommand</i>]</code>	Clears all the configuration for a specified command. If you only want to clear the configuration for a specific version of the command, you can enter a value for <i>level2configurationcommand</i> .
Example: <code>ciscoasa(config)# clear configure aaa</code>	For example, to clear the configuration for all aaa commands, enter the following command: <code>ciscoasa(config)# clear configure aaa</code> To clear the configuration for only aaa authentication commands, enter the following command: <code>ciscoasa(config)# clear configure aaa authentication</code>
<code>no <i>configurationcommand</i></code> <code>[<i>level2configurationcommand</i>] <i>qualifier</i></code>	Disables the specific parameters or options of a command. In this case, you use the no command to remove the specific configuration identified by <i>qualifier</i> .
Example: <code>ciscoasa(config)# no nat (inside) 1</code>	For example, to remove a specific nat command, enter enough of the command to identify it uniquely as follows: <code>ciscoasa(config)# no nat (inside) 1</code>

Command	Purpose
write erase Example: <pre>ciscoasa(config)# write erase</pre>	Erases the startup configuration.
clear configure all Example: <pre>ciscoasa(config)# clear configure all</pre>	Erases the running configuration. Note In multiple context mode, if you enter clear configure all from the system configuration, you also remove all contexts and stop them from running. The context configuration files are not erased, and remain in their original location.

Creating Text Configuration Files Offline

This guide describes how to use the CLI to configure the ASA; when you save commands, the changes are written to a text file. Instead of using the CLI, however, you can edit a text file directly on your PC and paste a configuration at the configuration mode command-line prompt in its entirety, or line by line. Alternatively, you can download a text file to the ASA internal flash memory. See [Chapter 42, “Managing Software and Configurations,”](#) for information on downloading the configuration file to the ASA.

In most cases, commands described in this guide are preceded by a CLI prompt. The prompt in the following example is “ciscoasa(config)#”:

```
ciscoasa(config)# context a
```

In the text configuration file you are not prompted to enter commands, so the prompt is omitted as follows:

```
context a
```

For additional information about formatting the file, see [Appendix 48, “Using the Command-Line Interface.”](#)

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
<p>clear local-host [<i>ip_address</i>] [all]</p> <p>Example: ciscoasa(config)# clear local-host all</p>	<p>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.</p> <p>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 <i>ip_address</i> argument.</p>
<p>clear conn [all] [protocol {tcp udp}] [address <i>src_ip</i>[-<i>src_ip</i>] [netmask <i>mask</i>]] [port <i>src_port</i>[-<i>src_port</i>]] [address <i>dest_ip</i>[-<i>dest_ip</i>] [netmask <i>mask</i>]] [port <i>dest_port</i>[-<i>dest_port</i>]]</p> <p>Example: ciscoasa(config)# clear conn all</p>	<p>This command terminates connections in any state. See the show conn command to view all current connections.</p> <p>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.</p>

Reloading the ASA

To reload the ASA, enter the following command:

Command	Purpose
<p>reload</p> <p>Example: ciscoasa (config)# reload</p>	<p>Reloads the ASA.</p> <p>Note In multiple context mode, you can only reload from the system execution space.</p>



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 4-1](#)
- [Information About Feature Licenses, page 4-23](#)
- [Guidelines and Limitations, page 4-33](#)
- [Configuring Licenses, page 4-35](#)
- [Monitoring Licenses, page 4-40](#)
- [Feature History for Licensing, page 4-50](#)

Supported Feature Licenses Per Model

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

- [Licenses Per Model, page 4-1](#)
- [License Notes, page 4-18](#)
- [VPN License and Feature Compatibility, page 4-23](#)

Licenses Per Model

This section lists the feature licenses available for each model:

- [ASA 5505, page 4-3](#)
- [ASA 5510, page 4-4](#)
- [ASA 5520, page 4-5](#)

- ASA 5540, page 4-6
- ASA 5550, page 4-7
- ASA 5580, page 4-8
- ASA 5512-X, page 4-9
- ASA 5515-X, page 4-10
- ASA 5525-X, page 4-11
- ASA 5545-X, page 4-12
- ASA 5555-X, page 4-13
- ASA 5585-X with SSP-10, page 4-14
- ASA 5585-X with SSP-20, page 4-15
- ASA 5585-X with SSP-40 and -60, page 4-16
- ASA Services Module, page 4-17

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 4-32 for a list of unsupported features.

For detailed information about licenses, see the [“License Notes”](#) section on page 4-18.

ASA 5505

Table 4-1 ASA 5505 License Features

Licenses	Description (Base License in Plain Text)				Description (Security Plus Lic. in Plain Text)			
Firewall Licenses								
Botnet Traffic Filter	Disabled	<i>Opt. Time-based lic: Available</i>			Disabled	<i>Opt. Time-based lic: Available</i>		
Firewall Conns, Concurrent	10,000				25,000			
GTP/GPRS	No support				No support			
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>			Disabled	<i>Optional license: Available</i>		
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional license: 24</i>			2	<i>Optional license: 24</i>		
VPN Licenses								
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>			Disabled	<i>Optional license: Available</i>		
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>			Disabled	<i>Optional license: Available</i>		
AnyConnect Essentials	Disabled	<i>Optional license: Available (25 sessions)</i>			Disabled	<i>Optional license: Available (25 sessions)</i>		
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>			Disabled	<i>Optional license: Available</i>		
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>	10	25	2	<i>Optional Permanent or Time-based licenses:</i>	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)	<i>Opt. lic.: Strong (3DES/AES)</i>			Base (DES)	<i>Opt. lic.: Strong (3DES/AES)</i>		
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	<i>Opt. licenses:</i>	50	<i>Unlimited</i>	10 ³	<i>Opt. licenses:</i>	50	<i>Unlimited</i>
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 4-2 ASA 5510 License Features

Licenses	Description (Base License in Plain Text)					Description (Security Plus Lic. in Plain Text)						
Firewall Licenses												
Botnet Traffic Filter	Disabled		<i>Optional Time-based license: Available</i>			Disabled		<i>Optional Time-based license: Available</i>				
Firewall Conns, Concurrent	50,000					130,000						
GTP/GPRS	No support					No support						
Intercompany Media Eng.	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>				2	<i>Optional licenses:</i>					
		24	50	100			24	50	100			
VPN Licenses												
Adv. Endpoint Assessment	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect for Cisco VPN Phone	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect Essentials	Disabled		<i>Optional license: Available (250 sessions)</i>			Disabled		<i>Optional license: Available (250 sessions)</i>				
AnyConnect for Mobile	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect Premium (sessions)	2	<i>Optional Perm. or Time-based lic.:</i>				2	<i>Optional Perm. or Time-based lic.:</i>					
		10	25	50	100	250		10	25	50	100	250
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>					<i>Optional Shared licenses: Participant or Server. For the Server:</i>						
	<i>500-50,000 in increments of 500</i>			<i>50,000-545,000 in increments of 1000</i>		<i>500-50,000 in increments of 500</i>			<i>50,000-545,000 in increments of 1000</i>			
Total VPN (sessions), combined all types	250					250						
Other VPN (sessions)	250					250						
VPN Load Balancing	No support					Supported						
General Licenses												
Encryption	Base (DES)		<i>Opt. lic.: Strong (3DES/AES)</i>			Base (DES)		<i>Opt. lic.: Strong (3DES/AES)</i>				
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	<i>Optional licenses:</i>		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 4-3 ASA 5520 License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled		<i>Optional Time-based license: Available</i>								
Firewall Conns, Concurrent	280,000										
GTP/GPRS	Disabled		<i>Optional license: Available</i>								
Intercompany Media Eng.	Disabled		<i>Optional license: Available</i>								
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>			24	50	100	250	500	750	1000
VPN Licenses											
Adv. Endpoint Assessment	Disabled		<i>Optional license: Available</i>								
AnyConnect for Cisco VPN Phone	Disabled		<i>Optional license: Available</i>								
AnyConnect Essentials	Disabled		<i>Optional license: Available (750 sessions)</i>								
AnyConnect for Mobile	Disabled		<i>Optional license: Available</i>								
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>									
		10	25	50	100	250	500	750			
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>										
	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)		<i>Optional license: Strong (3DES/AES)</i>								
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	764										
Security Contexts	2	<i>Optional licenses:</i>			5	10	20				
Clustering	No support										
VLANs, Maximum	150										

ASA 5540

Table 4-4 ASA 5540 License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled		<i>Optional Time-based license: Available</i>								
Firewall Conns, Concurrent	400,000										
GTP/GPRS	Disabled		<i>Optional license: Available</i>								
Intercompany Media Eng.	Disabled		<i>Optional license: Available</i>								
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>		24	50	100	250	500	750	1000	2000
VPN Licenses											
Adv. Endpoint Assessment	Disabled		<i>Optional license: Available</i>								
AnyConnect for Cisco VPN Phone	Disabled		<i>Optional license: Available</i>								
AnyConnect Essentials	Disabled		<i>Optional license: Available (2500 sessions)</i>								
AnyConnect for Mobile	Disabled		<i>Optional license: Available</i>								
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>									
		10	25	50	100	250	500	750	1000	2500	
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>										
	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)		<i>Optional license: Strong (3DES/AES)</i>								
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	964										
Security Contexts	2	<i>Optional licenses:</i>		5	10	20	50				
Clustering	No support										
VLANs, Maximum	200										

ASA 5550

Table 4-5 ASA 5550 License Features

Licenses	Description (Base License in Plain Text)									
Firewall Licenses										
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>								
Firewall Conns, Concurrent	650,000									
GTP/GPRS	Disabled	<i>Optional license: Available</i>								
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>								
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>								
	24	50	100	250	500	750	1000	2000	3000	
VPN Licenses										
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>								
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>								
AnyConnect Essentials	Disabled	<i>Optional license: Available (5000 sessions)</i>								
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>								
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>								
	10	25	50	100	250	500	750	1000	2500	5000
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>									
	<i>500-50,000 in increments of 500</i>					<i>50,000-545,000 in increments of 1000</i>				
Total VPN (sessions), combined all types	5000									
Other VPN (sessions)	5000									
VPN Load Balancing	Supported									
General Licenses										
Encryption	Base (DES)	<i>Optional license: Strong (3DES/AES)</i>								
Failover	Active/Standby or Active/Active									
Interfaces of all types, Max.	1764									
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100	
Clustering	No support									
VLANs, Maximum	400									

ASA 5580

Table 4-6 ASA 5580 License Features

Licenses	Description (Base License in Plain Text)											
Firewall Licenses												
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>										
Firewall Conns, Concurrent	5580-20: 2,000,000						5580-40: 4,000,000					
GTP/GPRS	Disabled	<i>Optional license: Available</i>										
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>										
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>										
	24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹	
VPN Licenses												
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>										
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>										
AnyConnect Essentials	Disabled	<i>Optional license: Available (10000 sessions)</i>										
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>										
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>										
	10	25	50	100	250	500	750	1000	2500	5000	10,000	
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>											
	<i>500-50,000 in increments of 500</i>						<i>50,000-545,000 in increments of 1000</i>					
Total VPN (sessions), combined all types	10,000											
Other VPN (sessions)	10,000											
VPN Load Balancing	Supported											
General Licenses												
Encryption	Base (DES)	<i>Optional license: Strong (3DES/AES)</i>										
Failover	Active/Standby or Active/Active											
Interfaces of all types, Max.	4612											
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100	250		
Clustering	Disabled	<i>Optional license: Available for 8 units</i>										
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 4-7 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		<i>Optional Time-based license: Available</i>			Disabled		<i>Optional Time-based license: Available</i>				
Firewall Conns, Concurrent	100,000					250,000						
GTP/GPRS	No support					Disabled		<i>Optional license: Available</i>				
Intercompany Media Eng.	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>				2	<i>Optional licenses:</i>					
		24	50	100	250	500		24	50	100	250	500
VPN Licenses												
Adv. Endpoint Assessment	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect for Cisco VPN Phone	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect Essentials	Disabled		<i>Optional license: Available (250 sessions)</i>			Disabled		<i>Optional license: Available (250 sessions)</i>				
AnyConnect for Mobile	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
AnyConnect Premium (sessions)	2	<i>Optional Perm. or Time-based lic.:</i>				2	<i>Optional Perm. or Time-based lic.:</i>					
		10	25	50	100	250		10	25	50	100	250
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>					<i>Optional Shared licenses: Participant or Server. For the Server:</i>						
	<i>500-50,000 in increments of 500</i>			<i>50,000-545,000 in increments of 1000</i>		<i>500-50,000 in increments of 500</i>			<i>50,000-545,000 in increments of 1000</i>			
Total VPN (sessions), combined all types	250					250						
Other VPN (sessions)	250					250						
VPN Load Balancing	No support					Supported						
General Licenses												
Encryption	Base (DES)		<i>Opt. lic.: Strong (3DES/AES)</i>			Base (DES)		<i>Opt. lic.: Strong (3DES/AES)</i>				
Failover	No support					Active/Standby or Active/Active						
Interfaces of all types, Max.	716					916						
Security Contexts	No support					2	<i>Optional licenses:</i>		5			
Clustering	No Support					Supported for 2 units						
IPS Module	Disabled		<i>Optional license: Available</i>			Disabled		<i>Optional license: Available</i>				
VLANs, Maximum	50					100						

ASA 5515-X

Table 4-8 ASA 5515-X License Features

Licenses	Description (Base License in Plain Text)								
Firewall Licenses									
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>							
Firewall Conns, Concurrent	250,000								
GTP/GPRS	Disabled	<i>Optional license: Available</i>							
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>							
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>			24	50	100	250	500
VPN Licenses									
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>							
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>							
AnyConnect Essentials	Disabled	<i>Optional license: Available (250 sessions)</i>							
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>							
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>							
		10	25	50	100	250			
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>								
	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)	<i>Optional license: Strong (3DES/AES)</i>							
Failover	Active/Standby or Active/Active								
Interfaces of all types, Max.	916								
Security Contexts	2	<i>Optional licenses:</i>			5				
Clustering	Supported for 2 units								
IPS Module	Disabled	<i>Optional license: Available</i>							
VLANs, Maximum	100								

ASA 5525-X

Table 4-9 ASA 5525-X License Features

Licenses	Description (Base License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>									
Firewall Conns, Concurrent	500,000										
GTP/GPRS	Disabled	<i>Optional license: Available</i>									
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>									
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>			24	50	100	250	500	750	1000
VPN Licenses											
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>									
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>									
AnyConnect Essentials	Disabled	<i>Optional license: Available (750 sessions)</i>									
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>									
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>									
		10	25	50	100	250	500	750			
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>										
	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)	<i>Optional license: Strong (3DES/AES)</i>									
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	1316										
Security Contexts	2	<i>Optional licenses:</i>			5	10	20				
Clustering	Supported for 2 units										
IPS Module	Disabled	<i>Optional license: Available</i>									
VLANs, Maximum	200										

ASA 5545-X

Table 4-10 ASA 5545-X License Features

Licenses	Description (Base License in Plain Text)											
Firewall Licenses												
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>										
Firewall Conns, Concurrent	750,000											
GTP/GPRS	Disabled	<i>Optional license: Available</i>										
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>										
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>			24	50	100	250	500	750	1000	2000
VPN Licenses												
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>										
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>										
AnyConnect Essentials	Disabled	<i>Optional license: Available (2500 sessions)</i>										
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>										
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>										
		10	25	50	100	250	500	750	1000	2500		
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>											
		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)	<i>Optional license: Strong (3DES/AES)</i>										
Failover	Active/Standby or Active/Active											
Interfaces of all types, Max.	1716											
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50				
Clustering	Supported for 2 units											
IPS Module	Disabled	<i>Optional license: Available</i>										
VLANs, Maximum	300											

ASA 5555-X

Table 4-11 ASA 5555-X License Features

Licenses	Description (Base License in Plain Text)									
Firewall Licenses										
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>								
Firewall Conns, Concurrent	1,000,000									
GTP/GPRS	Disabled	<i>Optional license: Available</i>								
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>								
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>								
	24	50	100	250	500	750	1000	2000	3000	
VPN Licenses										
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>								
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>								
AnyConnect Essentials	Disabled	<i>Optional license: Available (5000 sessions)</i>								
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>								
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>								
	10	25	50	100	250	500	750	1000	2500	5000
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>									
	<i>500-50,000 in increments of 500</i>					<i>50,000-545,000 in increments of 1000</i>				
Total VPN (sessions), combined all types	5000									
Other VPN (sessions)	5000									
VPN Load Balancing	Supported									
General Licenses										
Encryption	Base (DES)	<i>Optional license: Strong (3DES/AES)</i>								
Failover	Active/Standby or Active/Active									
Interfaces of all types, Max.	2516									
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100	
Clustering	Supported for 2 units									
IPS Module	Disabled	<i>Optional license: Available</i>								
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 4-12 ASA 5585-X with SSP-10 License Features

Licenses	Description (Base and Security Plus License in Plain Text)										
Firewall Licenses											
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>									
Firewall Conns, Concurrent	1,000,000										
GTP/GPRS	Disabled	<i>Optional license: Available</i>									
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>									
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>									
		24	50	100	250	500	750	1000	2000	3000	
VPN Licenses											
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>									
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>									
AnyConnect Essentials	Disabled	<i>Optional license: Available (5000 sessions)</i>									
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>									
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>									
		10	25	50	100	250	500	750	1000	2500	5000
		<i>Optional Shared licenses: Participant or Server. For the Server:</i>					500-50,000 in increments of 500				
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)	<i>Optional license: Strong (3DES/AES)</i>									
Failover	Active/Standby or Active/Active										
Interfaces of all types, Max.	4612										
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100		
Clustering	Disabled	<i>Optional license: Available for 8 units</i>									
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 4-13 ASA 5585-X with SSP-20 License Features

Licenses	Description (Base and Security Plus License in Plain Text)											
Firewall Licenses												
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>										
Firewall Conns, Concurrent	2,000,000											
GTP/GPRS	Disabled	<i>Optional license: Available</i>										
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>										
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>										
	24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹	
VPN Licenses												
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>										
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>										
AnyConnect Essentials	Disabled	<i>Optional license: Available (10,000 sessions)</i>										
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>										
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>										
	10	25	50	100	250	500	750	1000	2500	5000	10,000	
	<i>Optional Shared licenses: Participant or Server. For the Server:</i>											
	<i>500-50,000 in increments of 500</i>						<i>50,000-545,000 in increments of 1000</i>					
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)	<i>Optional license: Strong (3DES/AES)</i>										
Failover	Active/Standby or Active/Active											
Interfaces of all types, Max.	4612											
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100	250		
Clustering	Disabled	<i>Optional license: Available for 8 units</i>										
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 4-14 ASA 5585-X with SSP-40 and -60 License Features

Licenses	Description (Base License in Plain Text)												
Firewall Licenses													
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>											
Firewall Conns, Concurrent	5585-X with SSP-40: 4,000,000						5585-X with SSP-60: 10,000,000						
GTP/GPRS	Disabled	<i>Optional license: Available</i>											
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>											
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>											
		24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹	
VPN Licenses													
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>											
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>											
AnyConnect Essentials	Disabled	<i>Optional license: Available (10,000 sessions)</i>											
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>											
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>											
		10	25	50	100	250	500	750	1000	2500	5000	10,000	
		<i>Optional Shared licenses: Participant or Server. For the Server:</i>											
		<i>500-50,000 in increments of 500</i>						<i>50,000-545,000 in increments of 1000</i>					
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)	<i>Optional license: Strong (3DES/AES)</i>											
Failover	Active/Standby or Active/Active												
Interfaces of all types, Max.	4612												
Security Contexts	2	<i>Optional licenses:</i>			5	10	20	50	100	250			
Clustering	Disabled	<i>Optional license: Available for 8 units</i>											
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 4-15 ASASM License Features

Licenses	Description (Base License in Plain Text)											
Firewall Licenses												
Botnet Traffic Filter	Disabled	<i>Optional Time-based license: Available</i>										
Firewall Conns, Concurrent	10,000,000											
GTP/GPRS	Disabled	<i>Optional license: Available</i>										
Intercompany Media Eng.	Disabled	<i>Optional license: Available</i>										
UC Phone Proxy Sessions, Total UC Proxy Sessions	2	<i>Optional licenses:</i>										
		24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹
VPN Licenses												
Adv. Endpoint Assessment	Disabled	<i>Optional license: Available</i>										
AnyConnect for Cisco VPN Phone	Disabled	<i>Optional license: Available</i>										
AnyConnect Essentials	Disabled	<i>Optional license: Available (10,000 sessions)</i>										
AnyConnect for Mobile	Disabled	<i>Optional license: Available</i>										
AnyConnect Premium (sessions)	2	<i>Optional Permanent or Time-based licenses:</i>										
		10	25	50	100	250	500	750	1000	2500	5000	10,000
		<i>Optional Shared licenses: Participant or Server. For the Server:</i>										
		<i>500-50,000 in increments of 500</i>					<i>50,000-545,000 in increments of 1000</i>					
Total VPN (sessions), combined all types	10,000											
Other VPN (sessions)	10,000											
VPN Load Balancing	Supported											
General Licenses												
Encryption	Base (DES)	<i>Optional license: Strong (3DES/AES)</i>										
Failover	Active/Standby or Active/Active											
Security Contexts	2	<i>Optional licenses:</i>										
		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 4-16 includes common footnotes shared by multiple tables in the “Licenses Per Model” section on page 4-1.

Table 4-16 License Notes

License	Notes
AnyConnect Essentials	<p>AnyConnect Essentials sessions include the following VPN types:</p> <ul style="list-style-type: none"> • SSL VPN • IPsec remote access VPN using IKEv2 <p>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.</p> <p>Note With the AnyConnect Essentials license, VPN users can use a web browser to log in, and download and start (WebLaunch) the AnyConnect client.</p> <p>The AnyConnect client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium license.</p> <p>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.</p> <p>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.</p> <p>See also the “VPN License and Feature Compatibility” section on page 4-23.</p>
AnyConnect for Cisco VPN Phone	<p>In conjunction with an AnyConnect Premium license, this license enables access from hardware IP phones that have built in AnyConnect compatibility.</p>

Table 4-16 License Notes (continued)

License	Notes
AnyConnect for Mobile	<p>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.</p> <p>Mobile Posture Support</p> <p>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.</p> <ul style="list-style-type: none"> • AnyConnect Premium License Functionality <ul style="list-style-type: none"> – 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 <ul style="list-style-type: none"> – 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	<p>AnyConnect Premium sessions include the following VPN types:</p> <ul style="list-style-type: none"> • SSL VPN • Clientless SSL VPN • IPsec remote access VPN using IKEv2
AnyConnect Premium Shared	<p>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.</p>
Botnet Traffic Filter	<p>Requires a Strong Encryption (3DES/AES) License to download the dynamic database.</p>
Encryption	<p>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.</p>
Failover, Active/Active	<p>You cannot use Active/Active failover and VPN; if you want to use VPN, use Active/Standby failover.</p>

Table 4-16 License Notes (continued)

License	Notes
Intercompany Media Engine	<p>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.</p> <ul style="list-style-type: none"> • 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. <p>Note K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.</p> <p>You might also use SRTP encryption sessions for your connections:</p> <ul style="list-style-type: none"> • For a K8 license, SRTP sessions are limited to 250. • For a K9 license, there is no limit. <p>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.</p>
Interfaces of all types, Max.	<p>The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every interface command defined in the configuration counts against this limit. For example, both of the following interfaces count even if the GigabitEthernet 0/0 interface is defined as part of port-channel 1:</p> <pre>interface gigabitethernet 0/0</pre> <p>and</p> <pre>interface port-channel 1</pre>

Table 4-16 License Notes (continued)

License	Notes
IPS module	<p>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.</p> <p>See the following guidelines:</p> <ul style="list-style-type: none"> • 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.
Other VPN	<p>Other VPN sessions include the following VPN types:</p> <ul style="list-style-type: none"> • IPsec remote access VPN using IKEv1 • IPsec site-to-site VPN using IKEv1 • IPsec site-to-site VPN using IKEv2 <p>This license is included in the Base license.</p>
Total VPN (sessions), combined all types	<ul style="list-style-type: none"> • 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 4-16 License Notes (continued)

License	Notes
UC Phone Proxy sessions, Total UC Proxy Sessions	<p>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:</p> <ul style="list-style-type: none"> • Phone Proxy • Presence Federation Proxy • Encrypted Voice Inspection <p>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).</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>You might also use SRTP encryption sessions for your connections:</p> <ul style="list-style-type: none"> • For K8 licenses, SRTP sessions are limited to 250. • For K9 licenses, there is not limit. <p>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.</p>
VLANs, Maximum	<p>For an interface to count against the VLAN limit, you must assign a VLAN to it. For example:</p> <pre>interface gigabitethernet 0/0.100 vlan 100</pre>
VPN Load Balancing	VPN load balancing requires a Strong Encryption (3DES/AES) License.

VPN License and Feature Compatibility

Table 4-17 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 4-17 VPN License and Feature Compatibility

Supported with:	Enable one of the following licenses: ¹	
	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 **no anyconnect-essentials** command to enable the Premium license instead.
2. Mobile Posture support is different for the AnyConnect Essentials vs. the AnyConnect Premium license. See [Table 4-16 on page 4-18](#) 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 4-24](#)
- [Permanent License, page 4-24](#)
- [Time-Based Licenses, page 4-24](#)
- [Shared AnyConnect Premium Licenses, page 4-27](#)

- [Failover or ASA Cluster Licenses, page 4-30](#)
- [No Payload Encryption Models, page 4-32](#)
- [Licenses FAQ, page 4-33](#)

Preinstalled License

By default, your ASA ships with a license already installed. This license might be the Base License, to which you want to add more licenses, or it might already have all of your licenses installed, depending on what you ordered and what your vendor installed for you. See the [“Monitoring Licenses” section on page 4-40](#) section to determine which licenses you have installed.

Permanent License

You can have one permanent activation key installed. The permanent activation key includes all licensed features in a single key. If you also install time-based licenses, the ASA combines the permanent and time-based licenses into a running license. See the [“How Permanent and Time-Based Licenses Combine” section on page 4-25](#) 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 4-24](#)
- [How the Time-Based License Timer Works, page 4-25](#)
- [How Permanent and Time-Based Licenses Combine, page 4-25](#)
- [Stacking Time-Based Licenses, page 4-26](#)
- [Time-Based License Expiration, page 4-26](#)

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 4-18](#) 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 4-18 *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 4-40.

Stacking Time-Based Licenses

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

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

For example:

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

Similarly:

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

If the licenses are not identical (for example, a 1000-session AnyConnect Premium license vs. a 2500-session license), then the licenses are *not* combined. Because only one time-based license per feature can be active, only one of the licenses can be active. See the [“Activating or Deactivating Keys”](#) section on page 4-36 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 4-26 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 4-36.

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 4-27](#)
- [Communication Issues Between Participant and Server, page 4-28](#)
- [Information About the Shared Licensing Backup Server, page 4-28](#)
- [Failover and Shared Licenses, page 4-29](#)
- [Maximum Number of Participants, page 4-29](#)

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 4-29](#)
- [“Failover and Shared License Participants” section on page 4-29](#)

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.

The standby backup server shares the same operating limits as the primary backup server; if the standby unit becomes active, it continues counting down where the primary unit left off. See the [“Information About the Shared Licensing Backup Server” section on page 4-28](#) 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 4-30](#)
- [ASA Cluster License Requirements and Exceptions, page 4-30](#)
- [How Failover or ASA Cluster Licenses Combine, page 4-31](#)
- [Loss of Communication Between Failover or ASA Cluster Units, page 4-32](#)
- [Upgrading Failover Pairs, page 4-32](#)

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 4-40](#).

Loss of Communication Between Failover or ASA Cluster Units

If the units lose communication for more than 30 days, then each unit reverts to the license installed locally. During the 30-day grace period, the combined running license continues to be used by all units.

If you restore communication during the 30-day grace period, then for time-based licenses, the time elapsed is subtracted from the primary/master license; if the primary/master license becomes expired, only then does the secondary/slave license start to count down.

If you do not restore communication during the 30-day period, then for time-based licenses, time is subtracted from all unit licenses, if installed. They are treated as separate licenses and do not benefit from the combined license. The time elapsed includes the 30-day grace period.

For example:

1. You have a 52-week Botnet Traffic Filter license installed on two units. The combined running license allows a total duration of 104 weeks.
2. The units operate as a failover unit/ASA cluster for 10 weeks, leaving 94 weeks on the combined license (42 weeks on the primary/master, and 52 weeks on the secondary/slave).
3. If the units lose communication (for example the primary/master unit fails), the secondary/slave unit continues to use the combined license, and continues to count down from 94 weeks.
4. The time-based license behavior depends on when communication is restored:
 - Within 30 days—The time elapsed is subtracted from the primary/master unit license. In this case, communication is restored after 4 weeks. Therefore, 4 weeks are subtracted from the primary/master license leaving 90 weeks combined (38 weeks on the primary, and 52 weeks on the secondary).
 - After 30 days—The time elapsed is subtracted from both units. In this case, communication is restored after 6 weeks. Therefore, 6 weeks are subtracted from both the primary/master and secondary/slave licenses, leaving 84 weeks combined (36 weeks on the primary/master, and 46 weeks on the secondary/slave).

Upgrading Failover Pairs

Because failover pairs do not require the same license on both units, you can apply new licenses to each unit without any downtime. If you apply a permanent license that requires a reload (see [Table 4-19 on page 4-36](#)), 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 4-40](#)), VPN and Unified Communications licenses will not be listed.

Licenses FAQ

- Q.** Can I activate multiple time-based licenses, for example, AnyConnect Premium and Botnet Traffic Filter?
- A.** Yes. You can use one time-based license per feature at a time.
- Q.** Can I “stack” time-based licenses so that when the time limit runs out, it will automatically use the next license?
- A.** Yes. For identical licenses, the time limit is combined when you install multiple time-based licenses. For non-identical licenses (for example, a 1000-session AnyConnect Premium license and a 2500-session license), the ASA automatically activates the next time-based license it finds for the feature.
- Q.** Can I install a new permanent license while maintaining an active time-based license?
- A.** Yes. Activating a permanent license does not affect time-based licenses.
- Q.** For failover, can I use a shared licensing server as the primary unit, and the shared licensing backup server as the secondary unit?
- A.** No. The secondary unit has the same running license as the primary unit; in the case of the shared licensing server, they require a server license. The backup server requires a participant license. The backup server can be in a separate failover pair of two backup servers.
- Q.** Do I need to buy the same licenses for the secondary unit in a failover pair?
- A.** No. Starting with Version 8.3(1), you do not have to have matching licenses on both units. Typically, you buy a license only for the primary unit; the secondary unit inherits the primary license when it becomes active. In the case where you also have a separate license on the secondary unit (for example, if you purchased matching licenses for pre-8.3 software), the licenses are combined into a running failover cluster license, up to the model limits.
- Q.** Can I use a time-based or permanent AnyConnect Premium license in addition to a shared AnyConnect Premium license?
- A.** Yes. The shared license is used only after the sessions from the locally installed license (time-based or permanent) are used up. **Note:** On the shared licensing server, the permanent AnyConnect Premium license is not used; you can however use a time-based license at the same time as the shared licensing server license. In this case, the time-based license sessions are available for local AnyConnect Premium sessions only; they cannot be added to the shared licensing pool for use by participants.

Guidelines and Limitations

See the following guidelines for activation keys.

Context Mode Guidelines

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

Firewall Mode Guidelines

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

Failover Guidelines

- Shared licenses are not supported in Active/Active mode. See the [“Failover and Shared Licenses” section on page 4-29](#) for more information.
- See the [“Failover or ASA Cluster Licenses” section on page 4-30](#).

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 **no anyconnect-essentials** command.

Configuring Licenses

This section includes the following topics:

- [Obtaining an Activation Key, page 4-35](#)
- [Activating or Deactivating Keys, page 4-36](#)
- [Configuring a Shared License, page 4-37](#)

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 entering the following command.

```
ciscoasa# show version | grep Serial
```

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 4-19](#) lists the licenses that require reloading.

Table 4-19 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

	Command	Purpose
Step 1	<p>activation-key <i>key</i> [activate deactivate]</p> <p>Example: <pre>ciscoasa# activation-key 0xd11b3d48 0xa80a4c0a 0x48e0fd1c 0xb0443480 0x843fc490</pre></p>	<p>Applies an activation key to the ASA. The <i>key</i> 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.</p> <p>You can install one permanent key, and multiple time-based keys. If you enter a new permanent key, it overwrites the already installed one.</p> <p>The activate and deactivate keywords are available for time-based keys only. If you do not enter any value, activate is the default. The last time-based key that you activate for a given feature is the active one. To deactivate any active time-based key, enter the deactivate keyword. If you enter a key for the first time, and specify deactivate, then the key is installed on the ASA in an inactive state. See the “Time-Based Licenses” section on page 4-24 for more information.</p>
Step 2	<p>(Might be required.)</p> <p>reload</p> <p>Example: <pre>ciscoasa# reload</pre></p>	<p>Reloads the ASA. Some permanent licenses require you to reload the ASA after entering the new activation key. See Table 4-19 on page 4-36 for a list of licenses that need reloading. If you need to reload, you will see the following message:</p> <p>WARNING: The running activation key was not updated with the requested key. The flash activation key was updated with the requested key, and will become active after the next reload.</p>

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 4-27](#).

This section includes the following topics:

- [Configuring the Shared Licensing Server, page 4-37](#)
- [Configuring the Shared Licensing Backup Server \(Optional\), page 4-39](#)
- [Configuring the Shared Licensing Participant, page 4-39](#)

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

	Command	Purpose
Step 1	license-server secret <i>secret</i> Example: ciscoasa(config)# license-server secret farscape	Sets the shared secret, a string between 4 and 128 ASCII characters. Any participant with this secret can use the licensing server.
Step 2	(Optional) license-server refresh-interval <i>seconds</i> Example: ciscoasa(config)# license-server refresh-interval 100	Sets 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 3	(Optional) license-server port <i>port</i> Example: ciscoasa(config)# license-server port 40000	Sets 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) license-server backup <i>address backup-id serial_number [ha-backup-id ha_serial_number]</i> Example: ciscoasa(config)# license-server backup 10.1.1.2 backup-id JMX0916L0Z4 ha-backup-id JMX1378N0W3	Identifies the backup server IP address and serial number. If the backup server is part of a failover pair, identify the standby unit serial number as well. You can only identify 1 backup server and its optional standby unit.
Step 5	license-server enable <i>interface_name</i> Example: ciscoasa(config)# license-server enable inside	Enables this unit to be the shared licensing server. Specify the interface on which participants contact the server. You can repeat this command for as many interfaces as desired.

Examples

The following example sets the shared secret, changes the refresh interval and port, configures a backup server, and enables this unit as the shared licensing server on the inside interface and dmz interface:

```
ciscoasa(config)# license-server secret farscape
ciscoasa(config)# license-server refresh-interval 100
ciscoasa(config)# license-server port 40000
ciscoasa(config)# license-server backup 10.1.1.2 backup-id JMX0916L0Z4 ha-backup-id JMX1378N0W3
ciscoasa(config)# license-server enable inside
```

```
ciscoasa(config)# license-server enable dmz
```

What to Do Next

See the “[Configuring the Shared Licensing Backup Server \(Optional\)](#)” section on page 4-39, or the “[Configuring the Shared Licensing Participant](#)” section on page 4-39.

Configuring the Shared Licensing Backup Server (Optional)

This section enables a shared license participant to act as the backup server if the main server goes down.

Prerequisites

The backup server must have a shared licensing participant key.

Detailed Steps

	Command	Purpose
Step 1	<pre>license-server address address secret secret [port port]</pre> <p>Example: <pre>ciscoasa(config)# license-server address 10.1.1.1 secret farscape</pre></p>	Identifies the shared licensing server IP address and shared secret. If you changed the default port in the server configuration, set the port for the backup server to match.
Step 2	<pre>license-server backup enable interface_name</pre> <p>Example: <pre>ciscoasa(config)# license-server backup enable inside</pre></p>	Enables this unit to be the shared licensing backup server. Specify the interface on which participants contact the server. You can repeat this command for as many interfaces as desired.

Examples

The following example identifies the license server and shared secret, and enables this unit as the backup shared license server on the inside interface and dmz interface:

```
ciscoasa(config)# license-server address 10.1.1.1 secret farscape
ciscoasa(config)# license-server backup enable inside
ciscoasa(config)# license-server backup enable dmz
```

What to Do Next

See the “[Configuring the Shared Licensing Participant](#)” section on page 4-39.

Configuring the Shared Licensing Participant

This section configures a shared licensing participant to communicate with the shared licensing server.

Prerequisites

The participant must have a shared licensing participant key.

Detailed Steps

	Command	Purpose
Step 1	<pre>license-server address address secret secret [port port]</pre> <p>Example: ciscoasa(config)# license-server address 10.1.1.1 secret farscape</p>	Identifies the shared licensing server IP address and shared secret. If you changed the default port in the server configuration, set the port for the participant to match.
Step 2	<p>(Optional)</p> <pre>license-server backup address address</pre> <p>Example: ciscoasa(config)# license-server backup address 10.1.1.2</p>	If you configured a backup server, enter the backup server address.

Examples

The following example sets the license server IP address and shared secret, as well as the backup license server IP address:

```
ciscoasa(config)# license-server address 10.1.1.1 secret farscape
ciscoasa(config)# license-server backup address 10.1.1.2
```

Monitoring Licenses

This section includes the following topics:

- [Viewing Your Current License, page 4-40](#)
- [Monitoring the Shared License, page 4-49](#)

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 4-32](#) for more information.

Detailed Steps

Command	Purpose
<code>show activation-key [detail]</code>	This command shows the permanent license, active time-based licenses, and the running license, which is a combination of the permanent license and active time-based licenses. The detail keyword also shows inactive time-based licenses.
Example: <code>ciscoasa# show activation-key detail</code>	For failover units, this command also shows the “Failover cluster” license, which is the combined keys of the primary and secondary units.

Examples

Example 4-1 Standalone Unit Output for the show activation-key command

The following is sample output from the **show activation-key** command for a standalone unit that shows the running license (the combined permanent license and time-based licenses), as well as each active time-based license:

```
ciscoasa# show activation-key

Serial Number:  JMX1232L11M
Running Permanent Activation Key: 0xce06dc6b 0x8a7b5ab7 0xa1e21dd4 0xd2c4b8b8 0xc4594f9c
Running Timebased Activation Key: 0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285
Running Timebased Activation Key: 0xyadayad2 0xyadayad2 0xyadayad2 0xyadayad2 0xyadayad2
```

Licensed features for this platform:

```
Maximum Physical Interfaces      : Unlimited      perpetual
Maximum VLANs                   : 150           perpetual
Inside Hosts                    : Unlimited     perpetual
Failover                        : Active/Active perpetual
VPN-DES                          : Enabled       perpetual
VPN-3DES-AES                    : Enabled       perpetual
Security Contexts               : 10           perpetual
GTP/GPRS                        : Enabled       perpetual
AnyConnect Premium Peers        : 2            perpetual
AnyConnect Essentials           : Disabled     perpetual
Other VPN Peers                 : 750         perpetual
Total VPN Peers                 : 750         perpetual
Shared License                  : Enabled       perpetual
  Shared AnyConnect Premium Peers : 12000       perpetual
AnyConnect for Mobile           : Disabled     perpetual
AnyConnect for Cisco VPN Phone  : Disabled     perpetual
Advanced Endpoint Assessment    : Disabled     perpetual
UC Phone Proxy Sessions         : 12           62 days
Total UC Proxy Sessions         : 12           62 days
Botnet Traffic Filter           : Enabled      646 days
Intercompany Media Engine       : Disabled     perpetual
```

This platform has a Base license.

The flash permanent activation key is the SAME as the running permanent key.

```
Active Timebased Activation Key:
0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285
Botnet Traffic Filter           : Enabled      646 days

0xyadayad2 0xyadayad2 0xyadayad2 0xyadayad2 0xyadayad2
Total UC Proxy Sessions         : 10           62 days
```

Example 4-2 Standalone Unit Output for show activation-key detail

The following is sample output from the **show activation-key detail** command for a standalone unit that shows the running license (the combined permanent license and time-based licenses), as well as the permanent license and each installed time-based license (active and inactive):

```
ciscoasa# show activation-key detail

Serial Number: 88810093382
Running Permanent Activation Key: 0xce06dc6b 0x8a7b5ab7 0xa1e21dd4 0xd2c4b8b8 0xc4594f9c
Running Timebased Activation Key: 0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285

Licensed features for this platform:
Maximum Physical Interfaces      : 8                perpetual
VLANs                            : 20              DMZ Unrestricted
Dual ISPs                        : Enabled         perpetual
VLAN Trunk Ports                 : 8                perpetual
Inside Hosts                     : Unlimited      perpetual
Failover                         : Active/Standby perpetual
VPN-DES                          : Enabled         perpetual
VPN-3DES-AES                    : Enabled         perpetual
AnyConnect Premium Peers        : 2                perpetual
AnyConnect Essentials           : Disabled        perpetual
Other VPN Peers                  : 25              perpetual
Total VPN Peers                  : 25              perpetual
AnyConnect for Mobile           : Disabled        perpetual
AnyConnect for Cisco VPN Phone  : Disabled        perpetual
Advanced Endpoint Assessment     : Disabled        perpetual
UC Phone Proxy Sessions         : 2                perpetual
Total UC Proxy Sessions         : 2                perpetual
Botnet Traffic Filter           : Enabled         39 days
Intercompany Media Engine       : Disabled        perpetual

This platform has an ASA 5505 Security Plus license.

Running Permanent Activation Key: 0xce06dc6b 0x8a7b5ab7 0xa1e21dd4 0xd2c4b8b8 0xc4594f9c

Licensed features for this platform:
Maximum Physical Interfaces      : 8                perpetual
VLANs                            : 20              DMZ Unrestricted
Dual ISPs                        : Enabled         perpetual
VLAN Trunk Ports                 : 8                perpetual
Inside Hosts                     : Unlimited      perpetual
Failover                         : Active/Standby perpetual
VPN-DES                          : Enabled         perpetual
VPN-3DES-AES                    : Enabled         perpetual
AnyConnect Premium Peers        : 2                perpetual
AnyConnect Essentials           : Disabled        perpetual
Other VPN Peers                  : 25              perpetual
Total VPN Peers                  : 25              perpetual
AnyConnect for Mobile           : Disabled        perpetual
AnyConnect for Cisco VPN Phone  : Disabled        perpetual
Advanced Endpoint Assessment     : Disabled        perpetual
UC Phone Proxy Sessions         : 2                perpetual
Total UC Proxy Sessions         : 2                perpetual
Botnet Traffic Filter           : Enabled         39 days
Intercompany Media Engine       : Disabled        perpetual

The flash permanent activation key is the SAME as the running permanent key.

Active Timebased Activation Key:
```



```
0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285
Botnet Traffic Filter          : Enabled    39 days
```

```
Inactive Timebased Activation Key:
Oxyadayada3 Oxyadayada3 Oxyadayada3 Oxyadayada3 Oxyadayada3
AnyConnect Premium Peers      : 25      7 days
```

Example 4-3 Primary Unit Output in a Failover Pair for show activation-key detail

The following is sample output from the **show activation-key detail** command for the primary failover unit that shows:

- The primary unit license (the combined permanent license and time-based licenses).
- The “Failover Cluster” license, which is the combined licenses from the primary and secondary units. This is the license that is actually running on the ASA. The values in this license that reflect the combination of the primary and secondary licenses are in bold.
- The primary unit permanent license.
- The primary unit installed time-based licenses (active and inactive).

```
ciscoasa# show activation-key detail
```

```
Serial Number: P3000000171
Running Permanent Activation Key: 0xce06dc6b 0x8a7b5ab7 0xa1e21dd4 0xd2c4b8b8 0xc4594f9c
Running Timebased Activation Key: 0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285
```

```
Licensed features for this platform:
```

```
Maximum Physical Interfaces : Unlimited    perpetual
Maximum VLANs              : 150          perpetual
Inside Hosts               : Unlimited    perpetual
Failover                   : Active/Active perpetual
VPN-DES                    : Enabled      perpetual
VPN-3DES-AES               : Enabled      perpetual
Security Contexts          : 12          perpetual
GTP/GPRS                   : Enabled      perpetual
AnyConnect Premium Peers   : 2           perpetual
AnyConnect Essentials      : Disabled    perpetual
Other VPN Peers            : 750        perpetual
Total VPN Peers            : 750        perpetual
Shared License             : Disabled    perpetual
AnyConnect for Mobile      : Disabled    perpetual
AnyConnect for Cisco VPN Phone : Disabled    perpetual
Advanced Endpoint Assessment : Disabled    perpetual
UC Phone Proxy Sessions    : 2           perpetual
Total UC Proxy Sessions    : 2           perpetual
Botnet Traffic Filter      : Enabled     33 days
Intercompany Media Engine  : Disabled    perpetual
```

This platform has an ASA 5520 VPN Plus license.

```
Failover cluster licensed features for this platform:
```

```
Maximum Physical Interfaces : Unlimited    perpetual
Maximum VLANs              : 150          perpetual
Inside Hosts               : Unlimited    perpetual
Failover                   : Active/Active perpetual
VPN-DES                    : Enabled      perpetual
VPN-3DES-AES               : Enabled      perpetual
Security Contexts          : 12          perpetual
GTP/GPRS                   : Enabled      perpetual
AnyConnect Premium Peers : 4           perpetual
AnyConnect Essentials      : Disabled    perpetual
```

```

Other VPN Peers           : 750           perpetual
Total VPN Peers          : 750           perpetual
Shared License           : Disabled       perpetual
AnyConnect for Mobile    : Disabled       perpetual
AnyConnect for Cisco VPN Phone : Disabled       perpetual
Advanced Endpoint Assessment : Disabled       perpetual
UC Phone Proxy Sessions : 4           perpetual
Total UC Proxy Sessions : 4           perpetual
Botnet Traffic Filter     : Enabled       33 days
Intercompany Media Engine : Disabled       perpetual

```

This platform has an ASA 5520 VPN Plus license.

Running Permanent Activation Key: 0xce06dc6b 0x8a7b5ab7 0xa1e21dd4 0xd2c4b8b8 0xc4594f9c

Licensed features for this platform:

```

Maximum Physical Interfaces : Unlimited     perpetual
Maximum VLANs              : 150          perpetual
Inside Hosts               : Unlimited     perpetual
Failover                   : Active/Active perpetual
VPN-DES                    : Enabled       perpetual
VPN-3DES-AES               : Disabled     perpetual
Security Contexts          : 2            perpetual
GTP/GPRS                   : Disabled     perpetual
AnyConnect Premium Peers   : 2            perpetual
AnyConnect Essentials      : Disabled     perpetual
Other VPN Peers           : 750         perpetual
Total VPN Peers           : 750         perpetual
Shared License             : Disabled     perpetual
AnyConnect for Mobile      : Disabled     perpetual
AnyConnect for Cisco VPN Phone : Disabled     perpetual
Advanced Endpoint Assessment : Disabled     perpetual
UC Phone Proxy Sessions    : 2            perpetual
Total UC Proxy Sessions    : 2            perpetual
Botnet Traffic Filter      : Disabled     perpetual
Intercompany Media Engine  : Disabled     perpetual

```

The flash permanent activation key is the SAME as the running permanent key.

Active Timebased Activation Key:

```

0xa821d549 0x35725fe4 0xc918b97b 0xce0b987b 0x47c7c285
Botnet Traffic Filter : Enabled     33 days

```

Inactive Timebased Activation Key:

```

0xyadayad3 0xyadayad3 0xyadayad3 0xyadayad3 0xyadayad3
Security Contexts : 2            7 days
AnyConnect Premium Peers : 100         7 days

```

```

0xyadayad4 0xyadayad4 0xyadayad4 0xyadayad4 0xyadayad4
Total UC Proxy Sessions : 100         14 days

```

Example 4-4 Secondary Unit Output in a Failover Pair for show activation-key detail

The following is sample output from the **show activation-key detail** command for the secondary failover unit that shows:

- The secondary unit license (the combined permanent license and time-based licenses).
- The “Failover Cluster” license, which is the combined licenses from the primary and secondary units. This is the license that is actually running on the ASA. The values in this license that reflect the combination of the primary and secondary licenses are in bold.

- The secondary unit permanent license.
- The secondary installed time-based licenses (active and inactive). This unit does not have any time-based licenses, so none display in this sample output.

```
ciscoasa# show activation-key detail
```

```
Serial Number: P3000000011
Running Activation Key: Oxyadayad1 Oxyadayad1 Oxyadayad1 Oxyadayad1 Oxyadayad1
```

```
Licensed features for this platform:
```

```
Maximum Physical Interfaces : Unlimited perpetual
Maximum VLANs              : 150 perpetual
Inside Hosts               : Unlimited perpetual
Failover                   : Active/Active perpetual
VPN-DES                    : Enabled perpetual
VPN-3DES-AES               : Disabled perpetual
Security Contexts          : 2 perpetual
GTP/GPRS                   : Disabled perpetual
AnyConnect Premium Peers   : 2 perpetual
AnyConnect Essentials      : Disabled perpetual
Other VPN Peers            : 750 perpetual
Total VPN Peers            : 750 perpetual
Shared License              : Disabled perpetual
AnyConnect for Mobile      : Disabled perpetual
AnyConnect for Cisco VPN Phone : Disabled perpetual
Advanced Endpoint Assessment : Disabled perpetual
UC Phone Proxy Sessions    : 2 perpetual
Total UC Proxy Sessions    : 2 perpetual
Botnet Traffic Filter       : Disabled perpetual
Intercompany Media Engine   : Disabled perpetual
```

```
This platform has an ASA 5520 VPN Plus license.
```

```
Failover cluster licensed features for this platform:
```

```
Maximum Physical Interfaces : Unlimited perpetual
Maximum VLANs              : 150 perpetual
Inside Hosts               : Unlimited perpetual
Failover                   : Active/Active perpetual
VPN-DES                    : Enabled perpetual
VPN-3DES-AES               : Enabled perpetual
Security Contexts          : 10 perpetual
GTP/GPRS                   : Enabled perpetual
AnyConnect Premium Peers   : 4 perpetual
AnyConnect Essentials      : Disabled perpetual
Other VPN Peers            : 750 perpetual
Total VPN Peers            : 750 perpetual
Shared License              : Disabled perpetual
AnyConnect for Mobile      : Disabled perpetual
AnyConnect for Cisco VPN Phone : Disabled perpetual
Advanced Endpoint Assessment : Disabled perpetual
UC Phone Proxy Sessions    : 4 perpetual
Total UC Proxy Sessions    : 4 perpetual
Botnet Traffic Filter       : Enabled 33 days
Intercompany Media Engine   : Disabled perpetual
```

```
This platform has an ASA 5520 VPN Plus license.
```

```
Running Permanent Activation Key: Oxyadayad1 Oxyadayad1 Oxyadayad1 Oxyadayad1 Oxyadayad1
```

```
Licensed features for this platform:
```

```
Maximum Physical Interfaces : Unlimited perpetual
Maximum VLANs              : 150 perpetual
Inside Hosts               : Unlimited perpetual
```

```

Failover : Active/Active perpetual
VPN-DES : Enabled perpetual
VPN-3DES-AES : Disabled perpetual
Security Contexts : 2 perpetual
GTP/GPRS : Disabled perpetual
AnyConnect Premium Peers : 2 perpetual
AnyConnect Essentials : Disabled perpetual
Other VPN Peers : 750 perpetual
Total VPN Peers : 750 perpetual
Shared License : Disabled perpetual
AnyConnect for Mobile : Disabled perpetual
AnyConnect for Cisco VPN Phone : Disabled perpetual
Advanced Endpoint Assessment : Disabled perpetual
UC Phone Proxy Sessions : 2 perpetual
Total UC Proxy Sessions : 2 perpetual
Botnet Traffic Filter : Disabled perpetual
Intercompany Media Engine : Disabled perpetual

```

The flash permanent activation key is the SAME as the running permanent key.

Example 4-5 Primary Unit Output for the ASA Services Module in a Failover Pair for show activation-key

The following is sample output from the **show activation-key** command for the primary failover unit that shows:

- The primary unit license (the combined permanent license and time-based licenses).
- The “Failover Cluster” license, which is the combined licenses from the primary and secondary units. This is the license that is actually running on the ASA. The values in this license that reflect the combination of the primary and secondary licenses are in bold.
- The primary unit installed time-based licenses (active and inactive).

```
ciscoasa# show activation-key
```

```

erial Number: SAL144705BF
Running Permanent Activation Key: 0x4d1ed752 0xc8cfef37 0xf4c38198 0x93c04c28 0x4a1c049a
Running Timebased Activation Key: 0xbc07bbd7 0xb15591e0 0xed68c013 0xd79374ff 0x44f87880

```

```
Licensed features for this platform:
```

```

Maximum Interfaces : 1024 perpetual
Inside Hosts : Unlimited perpetual
Failover : Active/Active perpetual
DES : Enabled perpetual
3DES-AES : Enabled perpetual
Security Contexts : 25 perpetual
GTP/GPRS : Enabled perpetual
Botnet Traffic Filter : Enabled 330 days

```

This platform has an WS-SVC-ASA-SM1 No Payload Encryption license.

```
Failover cluster licensed features for this platform:
```

```

Maximum Interfaces : 1024 perpetual
Inside Hosts : Unlimited perpetual
Failover : Active/Active perpetual
DES : Enabled perpetual
3DES-AES : Enabled perpetual
Security Contexts : 50 perpetual
GTP/GPRS : Enabled perpetual
Botnet Traffic Filter : Enabled 330 days

```

This platform has an WS-SVC-ASA-SM1 No Payload Encryption license.

The flash permanent activation key is the SAME as the running permanent key.

```
Active Timebased Activation Key:
0xbc07bbd7 0xb15591e0 0xed68c013 0xd79374ff 0x44f87880
Botnet Traffic Filter           : Enabled    330 days
```

Example 4-6 Secondary Unit Output for the ASA Services Module in a Failover Pair for show activation-key

The following is sample output from the **show activation-key** command for the secondary failover unit that shows:

- The secondary unit license (the combined permanent license and time-based licenses).
- The “Failover Cluster” license, which is the combined licenses from the primary and secondary units. This is the license that is actually running on the ASA. The values in this license that reflect the combination of the primary and secondary licenses are in bold.
- The secondary installed time-based licenses (active and inactive). This unit does not have any time-based licenses, so none display in this sample output.

```
ciscoasa# show activation-key detail

Serial Number: SAD143502E3
Running Permanent Activation Key: 0xf404c46a 0xb8e5bd84 0x28c1b900 0x92eca09c 0x4e2a0683

Licensed features for this platform:
Maximum Interfaces           : 1024           perpetual
Inside Hosts                 : Unlimited        perpetual
Failover                     : Active/Active    perpetual
DES                           : Enabled          perpetual
3DES-AES                     : Enabled          perpetual
Security Contexts           : 25               perpetual
GTP/GPRS                     : Disabled         perpetual
Botnet Traffic Filter       : Disabled         perpetual

This platform has an WS-SVC-ASA-SM1 No Payload Encryption license.

Failover cluster licensed features for this platform:
Maximum Interfaces           : 1024           perpetual
Inside Hosts                 : Unlimited        perpetual
Failover                     : Active/Active    perpetual
DES                           : Enabled          perpetual
3DES-AES                     : Enabled          perpetual
Security Contexts           : 50             perpetual
GTP/GPRS                     : Enabled        perpetual
Botnet Traffic Filter       : Enabled        330 days
```

This platform has an WS-SVC-ASA-SM1 No Payload Encryption license.

The flash permanent activation key is the SAME as the running permanent key.

Example 4-7 Output in a Cluster for show activation-key

```
ciscoasa# show activation-key
Serial Number: JMX1504L2TD
Running Permanent Activation Key: 0x4a3eea7b 0x54b9f61a 0x4143a90c 0xe5849088 0x4412d4a9
```

```
Licensed features for this platform:
Maximum Physical Interfaces : Unlimited perpetual
Maximum VLANs : 100 perpetual
Inside Hosts : Unlimited perpetual
Failover : Active/Active perpetual
Encryption-DES : Enabled perpetual
Encryption-3DES-AES : Enabled perpetual
Security Contexts : 2 perpetual
GTP/GPRS : Disabled perpetual
AnyConnect Premium Peers : 2 perpetual
AnyConnect Essentials : Disabled perpetual
Other VPN Peers : 250 perpetual
Total VPN Peers : 250 perpetual
Shared License : Disabled perpetual
AnyConnect for Mobile : Disabled perpetual
AnyConnect for Cisco VPN Phone : Disabled perpetual
Advanced Endpoint Assessment : Disabled perpetual
UC Phone Proxy Sessions : 2 perpetual
Total UC Proxy Sessions : 2 perpetual
Botnet Traffic Filter : Disabled perpetual
Intercompany Media Engine : Disabled perpetual
Cluster : Enabled perpetual
```

This platform has an ASA 5585-X base license.

```
Failover cluster licensed features for this platform:
Maximum Physical Interfaces : Unlimited perpetual
Maximum VLANs : 100 perpetual
Inside Hosts : Unlimited perpetual
Failover : Active/Active perpetual
Encryption-DES : Enabled perpetual
Encryption-3DES-AES : Enabled perpetual
Security Contexts : 4 perpetual
GTP/GPRS : Disabled perpetual
AnyConnect Premium Peers : 4 perpetual
AnyConnect Essentials : Disabled perpetual
Other VPN Peers : 250 perpetual
Total VPN Peers : 250 perpetual
Shared License : Disabled perpetual
AnyConnect for Mobile : Disabled perpetual
AnyConnect for Cisco VPN Phone : Disabled perpetual
Advanced Endpoint Assessment : Disabled perpetual
UC Phone Proxy Sessions : 4 perpetual
Total UC Proxy Sessions : 4 perpetual
Botnet Traffic Filter : Disabled perpetual
Intercompany Media Engine : Disabled perpetual
Cluster : Enabled perpetual
```

This platform has an ASA 5585-X base license.

The flash permanent activation key is the SAME as the running permanent key.

Monitoring the Shared License

To monitor the shared license, enter one of the following commands.

Command	Purpose
<code>show shared license [detail client [hostname] backup]</code>	Shows shared license statistics. Optional keywords are available only for the licensing server: the detail keyword shows statistics per participant. To limit the display to one participant, use the client keyword. The backup keyword shows information about the backup server. To clear the shared license statistics, enter the clear shared license command.
<code>show activation-key</code>	Shows the licenses installed on the ASA. The show version command also shows license information.
<code>show vpn-sessiondb</code>	Shows license information about VPN sessions.

Examples

The following is sample output from the **show shared license** command on the license participant:

```
ciscoasa> show shared license
Primary License Server : 10.3.32.20
  Version               : 1
  Status                : Inactive

Shared license utilization:
SSLVPN:
  Total for network    :    5000
  Available            :    5000
  Utilized             :         0
This device:
  Platform limit      :    250
  Current usage       :         0
  High usage          :         0
Messages Tx/Rx/Error:
  Registration        : 0 / 0 / 0
  Get                 : 0 / 0 / 0
  Release             : 0 / 0 / 0
  Transfer            : 0 / 0 / 0
```

The following is sample output from the **show shared license detail** command on the license server:

```
ciscoasa> show shared license detail
Backup License Server Info:

Device ID             : ABCD
Address               : 10.1.1.2
Registered            : NO
HA peer ID           : EFGH
Registered            : NO
Messages Tx/Rx/Error:
  Hello               : 0 / 0 / 0
  Sync                : 0 / 0 / 0
  Update              : 0 / 0 / 0

Shared license utilization:
SSLVPN:
  Total for network    :    500
```

```

Available      :      500
Utilized      :          0
This device:
Platform limit :      250
Current usage  :          0
High usage     :          0
Messages Tx/Rx/Error:
Registration   : 0 / 0 / 0
Get           : 0 / 0 / 0
Release       : 0 / 0 / 0
Transfer      : 0 / 0 / 0

```

Client Info:

```

Hostname       : 5540-A
Device ID     : XXXXXXXXXXXX
SSLVPN:
Current usage  : 0
High          : 0
Messages Tx/Rx/Error:
Registration   : 1 / 1 / 0
Get           : 0 / 0 / 0
Release       : 0 / 0 / 0
Transfer      : 0 / 0 / 0
...

```

Feature History for Licensing

Table 4-20 lists each feature change and the platform release in which it was implemented.

Table 4-20 Feature History for Licensing

Feature Name	Platform Releases	Feature Information
Increased Connections and VLANs	7.0(5)	Increased the following limits: <ul style="list-style-type: none"> 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.

Table 4-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
Increased VLANs	7.2(2)	<p>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.</p> <p>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).</p>
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	<p>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.</p> <p>Note The interface names remain Ethernet 0/0 and Ethernet 0/1.</p> <p>Use the speed command to change the speed on the interface and use the show interface command to see what speed is currently configured for each interface.</p>
Advanced Endpoint Assessment License	8.0(2)	<p>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).</p> <p>With an Advanced Endpoint Assessment License, you can enhance Host Scan by configuring an attempt to update noncompliant computers to meet version requirements.</p> <p>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.</p>
VPN Load Balancing for the ASA 5510	8.0(2)	<p>VPN load balancing is now supported on the ASA 5510 Security Plus license.</p>

Table 4-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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)	<p>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.</p> <p>This feature is not available in Version 8.1.</p>
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)	<p>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.</p> <p>Note With the AnyConnect Essentials license, VPN users can use a Web browser to log in, and download and start (WebLaunch) the AnyConnect client.</p> <p>The AnyConnect client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium license.</p> <p>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.</p> <p>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.</p>
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.

Table 4-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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).
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 commands: show activation-key and show version .
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 <i>stack</i> 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. We modified the following commands: show activation-key and show version .
Discrete activation and deactivation of time-based licenses.	8.3(1)	You can now activate or deactivate time-based licenses using a command. We modified the following commands: activation-key [activate deactivate] .

Table 4-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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.
Increased connections for the ASA 5580 and 5585-X	8.4(1)	We increased the firewall connection limits: <ul style="list-style-type: none"> • 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. Note The following limitation exists in our support for IKEv2 on the ASA: We currently do not support duplicate security associations. 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.

Table 4-20 Feature History for Licensing (continued)

Feature Name	Platform Releases	Feature Information
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.
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 5-1](#)
- [Licensing Requirements for the Firewall Mode, page 5-7](#)
- [Default Settings, page 5-7](#)
- [Guidelines and Limitations, page 5-8](#)
- [Setting the Firewall Mode, page 5-9](#)
- [Configuring ARP Inspection for the Transparent Firewall, page 5-10](#)
- [Customizing the MAC Address Table for the Transparent Firewall, page 5-12](#)
- [Monitoring the Transparent Firewall, page 5-13](#)
- [Firewall Mode Examples, page 5-14](#)
- [Feature History for the Firewall Mode, page 5-25](#)

Information About the Firewall Mode

- [Information About Routed Firewall Mode, page 5-1](#)
- [Information About Transparent Firewall Mode, page 5-2](#)

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

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 5-1 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 5-1 Transparent Firewall Network



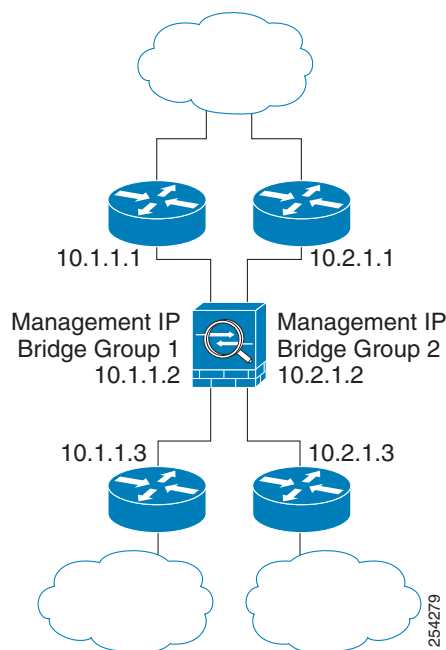
92411

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 5-2 shows two networks connected to the ASA, which has two bridge groups.

Figure 5-2 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 5-4.

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 9-2.

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 19, “Adding an Extended Access Control List,”](#) 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, 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 7-14 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 12-5](#) 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” section on page 5-9](#) 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. See the [“Configuring the Images and Startup Configuration to Use” section on page 42-21](#) for information about downloading text files.

Unsupported Features in Transparent Mode

Table 5-1 lists the features are not supported in transparent mode.

Table 5-1 *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 DHCP 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



Note

This section describes how to change the firewall mode. 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 5-8 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. See the [“Backing Up Configurations or Other Files”](#) section on page 42-25.
- 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 3-7 or [“Customizing ASDM Access \(ASA 5510 and Higher\)”](#) section on page 3-10.

Command	Purpose
<code>firewall transparent</code>	Sets the firewall mode to transparent. To change the mode to routed, enter the no firewall transparent command.
Example: <pre>ciscoasa(config)# firewall transparent</pre>	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 5-10](#)
- [Adding a Static ARP Entry, page 5-10](#)
- [Enabling ARP Inspection, page 5-11](#)

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 5-10. 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 5-11.
-

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

Command	Purpose
<code>arp interface_name ip_address mac_address</code>	Adds a static ARP entry.
Example: <pre>ciscoasa(config)# arp outside 10.1.1.1 0009.7cbe.2100</pre>	

Examples

For example, to allow ARP responses from the router at 10.1.1.1 with the MAC address 0009.7cbe.2100 on the outside interface, enter the following command:

```
ciscoasa(config)# arp outside 10.1.1.1 0009.7cbe.2100
```

What to Do Next

Enable ARP inspection according to the [“Enabling ARP Inspection”](#) section on page 5-11.

Enabling ARP Inspection

This section describes how to enable ARP inspection.

Detailed Steps

Command	Purpose
<code>arp-inspection interface_name enable</code> [<code>flood</code> <code>no-flood</code>]	Enables ARP inspection. The flood keyword forwards non-matching ARP packets out all interfaces, and no-flood drops non-matching packets.
Example: <pre>ciscoasa(config)# arp-inspection outside enable no-flood</pre>	Note The default setting is to flood non-matching packets. To restrict ARP through the ASA to only static entries, then set this command to no-flood .

Examples

For example, to enable ARP inspection on the outside interface, and to drop all non-matching ARP packets, enter the following command:

```
ciscoasa(config)# arp-inspection outside enable no-flood
```

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 5-12](#)
- [Setting the MAC Address Timeout, page 5-12](#)
- [Disabling MAC Address Learning, page 5-13](#)

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 5-10), a static MAC address entry is automatically added to the MAC address table.

To add a static MAC address to the MAC address table, enter the following command:

Command	Purpose
<pre>mac-address-table static interface_name mac_address</pre> <p>Example: <pre>ciscoasa(config)# mac-address-table static inside 0009.7cbe.2100</pre></p>	<p>Adds a static MAC address entry.</p> <p>The <i>interface_name</i> is the source interface.</p>

Setting the MAC Address Timeout

The default timeout value for dynamic MAC address table entries is 5 minutes, but you can change the timeout. To change the timeout, enter the following command:

Command	Purpose
<pre>mac-address-table aging-time timeout_value</pre> <p>Example: <pre>ciscoasa(config)# mac-address-table aging-time 10</pre></p>	<p>Sets the MAC address entry timeout.</p> <p>The <i>timeout_value</i> (in minutes) is between 5 and 720 (12 hours). 5 minutes is the default.</p>

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, enter the following command:

Command	Purpose
<code>mac-learn interface_name disable</code>	Disables MAC address learning.
Example: <code>ciscoasa(config)# mac-learn inside disable</code>	The no form of this command reenables MAC address learning. The clear configure mac-learn command reenables MAC address learning on all interfaces.

Monitoring the Transparent Firewall

- [Monitoring ARP Inspection, page 5-13](#)
- [Monitoring the MAC Address Table, page 5-13](#)

Monitoring ARP Inspection

To monitor ARP inspection, perform the following task:

Command	Purpose
<code>show arp-inspection</code>	Shows the current settings for ARP inspection on all interfaces.

Monitoring the MAC Address Table

You can view the entire MAC address table (including static and dynamic entries for both interfaces), or you can view the MAC address table for an interface. To view the MAC address table, enter the following command:

Command	Purpose
<code>show mac-address-table [interface_name]</code>	Shows the MAC address table.

Examples

The following is sample output from the `show mac-address-table` command that shows the entire table:

```
ciscoasa# show mac-address-table
interface      mac address      type      Time Left
-----
outside        0009.7cbe.2100   static    -
inside         0010.7cbe.6101   static    -
inside         0009.7cbe.5101   dynamic   10
```

The following is sample output from the **show mac-address-table** command that shows the table for the inside interface:

```
ciscoasa# show mac-address-table inside
interface      mac address      type      Time Left
-----
inside         0010.7cbe.6101   static    -
inside         0009.7cbe.5101   dynamic   10
```

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 5-14](#)
- [How Data Moves Through the Transparent Firewall, page 5-20](#)

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 5-15](#)
- [An Outside User Visits a Web Server on the DMZ, page 5-16](#)
- [An Inside User Visits a Web Server on the DMZ, page 5-17](#)
- [An Outside User Attempts to Access an Inside Host, page 5-17](#)
- [A DMZ User Attempts to Access an Inside Host, page 5-19](#)

An Inside User Visits a Web Server

Figure 5-3 shows an inside user accessing an outside web server.

Figure 5-3 Inside to Outside



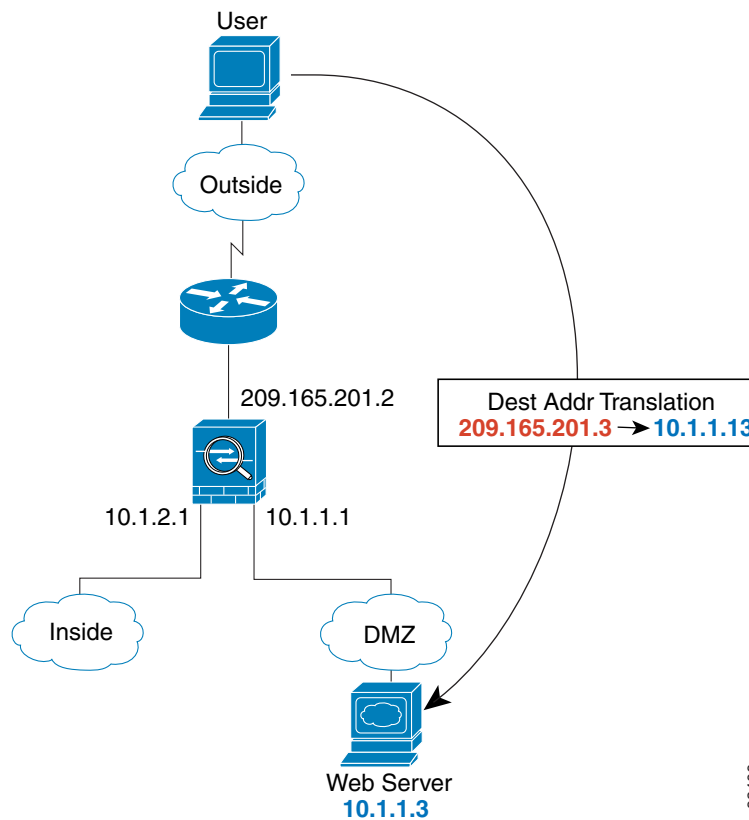
The following steps describe how data moves through the ASA (see Figure 5-3):

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 5-4 shows an outside user accessing the DMZ web server.

Figure 5-4 Outside to DMZ




The following steps describe how data moves through the ASA (see Figure 5-4):

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 5-5 shows an inside user accessing the DMZ web server.

Figure 5-5 Inside to DMZ

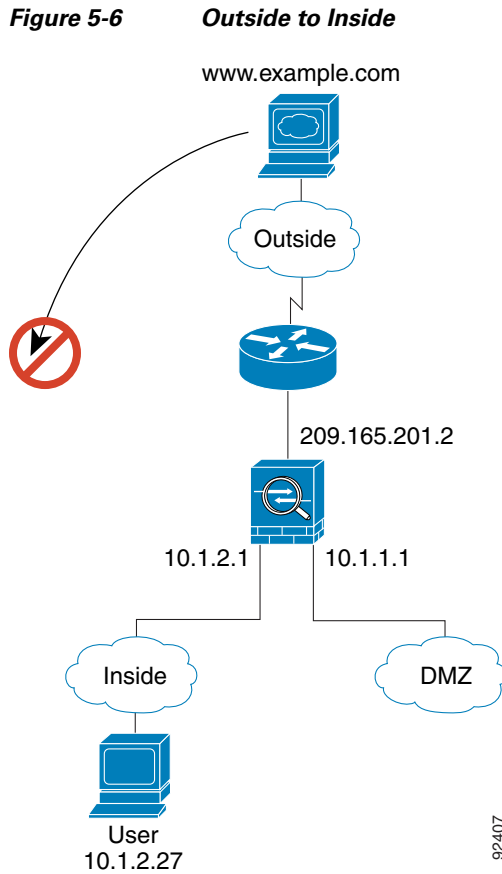


The following steps describe how data moves through the ASA (see Figure 5-5):

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 5-6 shows an outside user attempting to access the inside network.



92407

The following steps describe how data moves through the ASA (see [Figure 5-6](#)):

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.

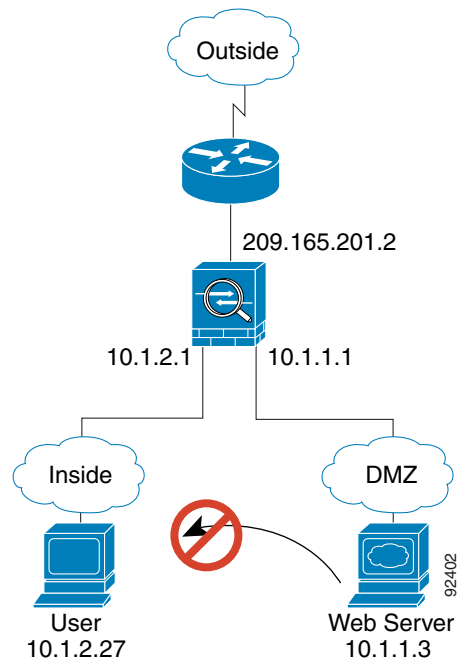
2. The ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according to the security policy (access lists, filters, AAA).
3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

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 5-7 shows a user in the DMZ attempting to access the inside network.

Figure 5-7 DMZ to Inside



The following steps describe how data moves through the ASA (see Figure 5-7):

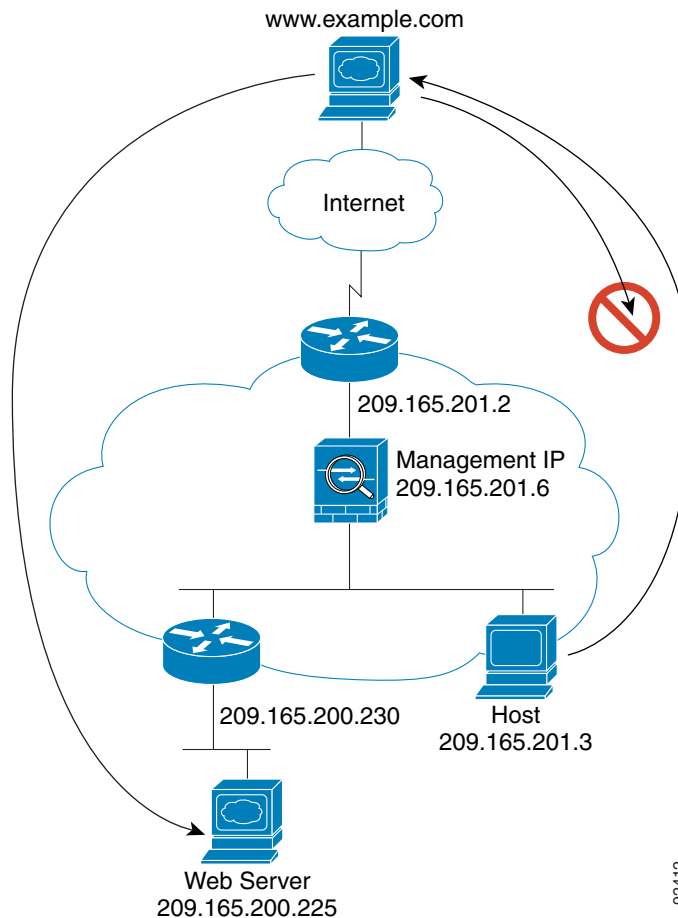
1. A user on the DMZ network attempts to reach an inside host. Because the DMZ does not have to route the traffic on the Internet, the private addressing scheme does not prevent routing.
2. The ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according 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 5-8 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 5-8 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 5-21](#)
- [An Inside User Visits a Web Server Using NAT, page 5-22](#)
- [An Outside User Visits a Web Server on the Inside Network, page 5-23](#)
- [An Outside User Attempts to Access an Inside Host, page 5-24](#)

An Inside User Visits a Web Server

Figure 5-9 shows an inside user accessing an outside web server.

Figure 5-9 *Inside to Outside*



The following steps describe how data moves through the ASA (see Figure 5-9):

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.

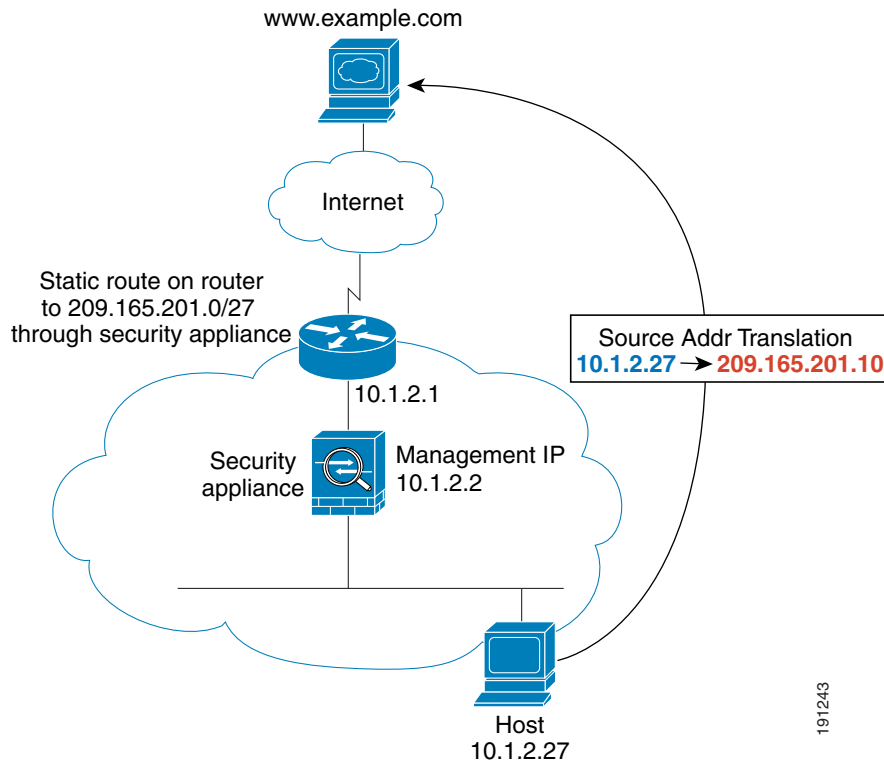
If the destination MAC address is not in the ASA table, the ASA attempts to discover the MAC address by sending an ARP request or a ping. The first packet is dropped.

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 5-10 shows an inside user accessing an outside web server.

Figure 5-10 Inside to Outside with NAT



The following steps describe how data moves through the ASA (see Figure 5-10):

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 5-11 shows an outside user accessing the inside web server.

Figure 5-11 *Outside to Inside*



The following steps describe how data moves through the ASA (see Figure 5-11):

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.

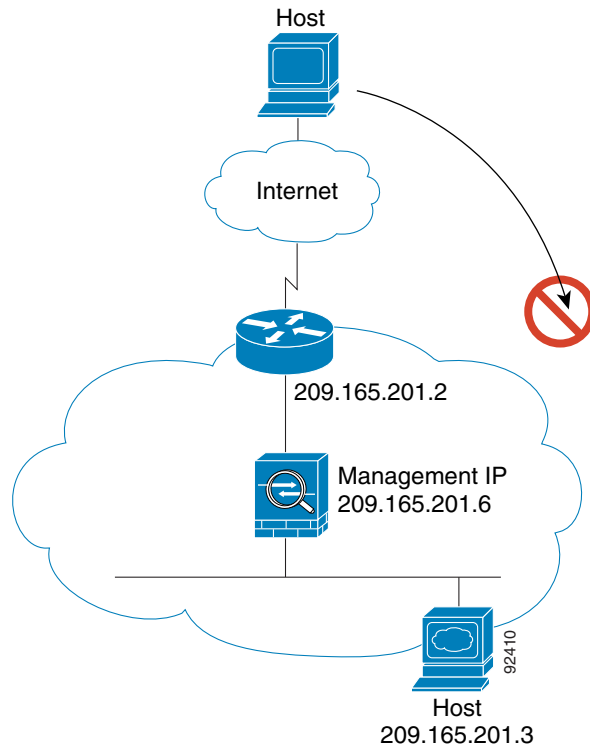
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.
If the destination MAC address is not in the ASA table, the ASA attempts 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 5-12 shows an outside user attempting to access a host on the inside network.

Figure 5-12 *Outside to Inside*



The following steps describe how data moves through the ASA (see Figure 5-12):

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 5-2 lists each feature change and the platform release in which it was implemented.

Table 5-2 Feature History for Firewall Mode

Feature Name	Platform Releases	Feature Information
Transparent Firewall Mode	7.0(1)	<p>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.</p> <p>We introduced the following commands: firewall transparent, show firewall.</p>
ARP inspection	7.0(1)	<p>ARP inspection compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table.</p> <p>We introduced the following commands: arp, arp-inspection, and show arp-inspection.</p>
MAC address table	7.0(1)	<p>Transparent firewall mode uses a MAC address table.</p> <p>We introduced the following commands: mac-address-table static, mac-address-table aging-time, mac-learn disable, and show mac-address-table.</p>
Transparent firewall bridge groups	8.4(1)	<p>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.</p> <p>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.</p> <p>We introduced the following commands: interface bvi, bridge-group, show bridge-group.</p>

Table 5-2 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)	<p>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.</p> <p>You may want to use this feature if you use:</p> <ul style="list-style-type: none"> • Secondary subnets. • Proxy ARP on adjacent routes for traffic forwarding. <p>We introduced the following command: arp permit-nonconnected.</p>
Mixed firewall mode support in multiple context mode	8.5(1)/9.0(1)	<p>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.</p> <p>We modified the following command: firewall transparent.</p>



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 6-1](#)
- [Licensing Requirements for Multiple Context Mode, page 6-13](#)
- [Guidelines and Limitations, page 6-14](#)
- [Default Settings, page 6-15](#)
- [Configuring Multiple Contexts, page 6-15](#)
- [Changing Between Contexts and the System Execution Space, page 6-24](#)
- [Managing Security Contexts, page 6-25](#)
- [Monitoring Security Contexts, page 6-28](#)
- [Configuration Examples for Multiple Context Mode, page 6-39](#)
- [Feature History for Multiple Context Mode, page 6-40](#)

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 6-14](#).

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

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

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 6-2](#)
- [System Configuration, page 6-2](#)
- [Admin Context Configuration, page 6-2](#)

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 6-3](#)
- [Classification Examples, page 6-4](#)

**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 6-3](#)
- [Unique MAC Addresses, page 6-3](#)
- [NAT Configuration, page 6-3](#)

**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.

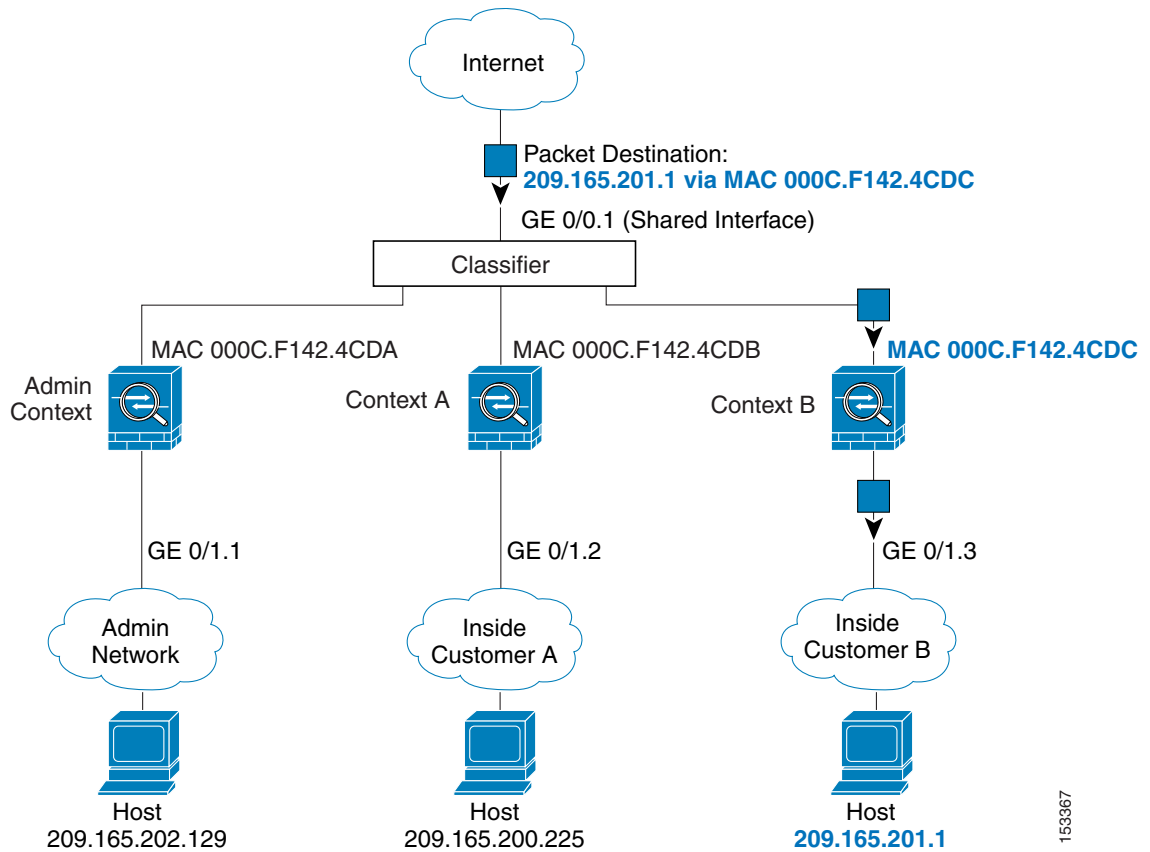
NAT Configuration

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 6-1 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 6-1 Packet Classification with a Shared Interface Using MAC Addresses



153367


Note that all new incoming traffic must be classified, even from inside networks. [Figure 6-2](#) 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 6-2 Incoming Traffic from Inside Networks



For transparent firewalls, you must use unique interfaces. [Figure 6-3](#) 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 6-3 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 6-4 shows a gateway context with two contexts behind the gateway.

Figure 6-4 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 6-7](#)
- [Context Administrator Access, page 6-8](#)

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 41, “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 41, “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 6-8](#)
- [Resource Limits, page 6-8](#)
- [Default Class, page 6-9](#)
- [Using Oversubscribed Resources, page 6-10](#)
- [Using Unlimited Resources, page 6-11](#)

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.

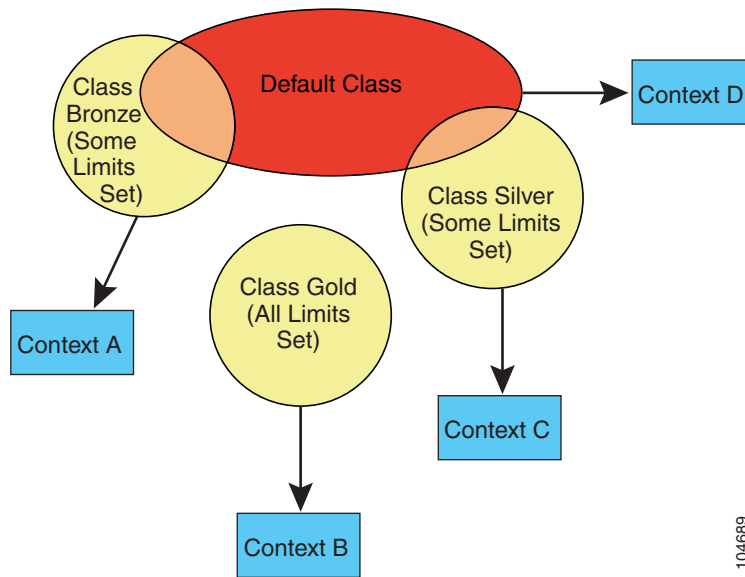
If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 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 6-5 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 6-5 Resource Classes

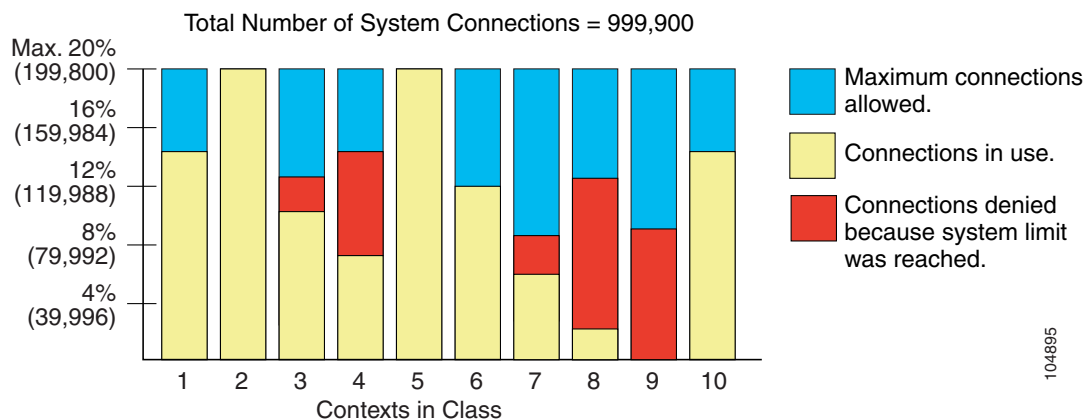


104689

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 6-6.)

Figure 6-6 Resource Oversubscription

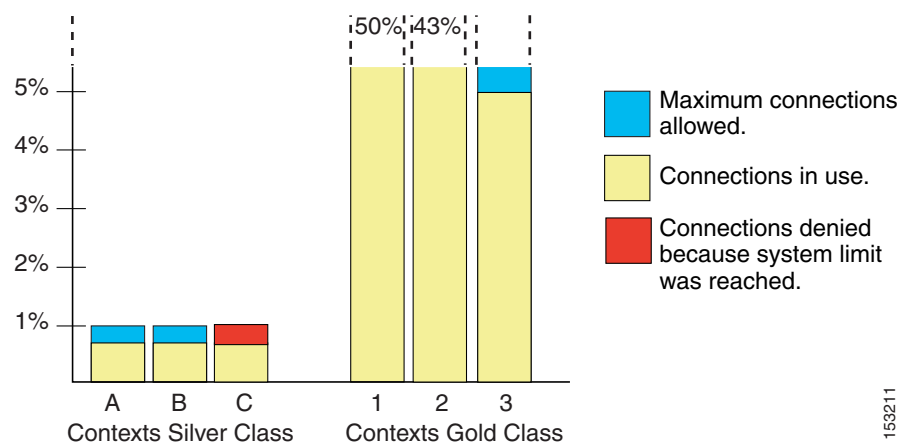


104895

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 6-7](#).) Setting unlimited access is similar to oversubscribing the ASA, except that you have less control over how much you oversubscribe the system.

Figure 6-7 Unlimited Resources



153211

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 6-24.

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 6-3 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 11-10 to manually set the MAC address.

This section includes the following topics:

- [Default MAC Address, page 6-12](#)
- [Interaction with Manual MAC Addresses, page 6-12](#)
- [Failover MAC Addresses, page 6-12](#)
- [MAC Address Format, page 6-12](#)

Default MAC Address

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 6-12](#).



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 6-12](#) section for more information.

MAC Address Format

The ASA generates the MAC address using the following format:

A2xx.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:

A24D.00zz.zzzz

For a prefix of 1009 (03F1), the MAC address is:

A2F1.03zz.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	<ul style="list-style-type: none"> Base License: No support. Security Plus License: 2 contexts. <p><i>Optional license: 5 contexts.</i></p>
ASA 5520	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, or 20 contexts.</i></p>
ASA 5540	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, or 50 contexts.</i></p>
ASA 5550	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, 50, or 100 contexts.</i></p>
ASA 5580	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.</i></p>
ASA 5512-X	<ul style="list-style-type: none"> Base License: No support. Security Plus License: 2 contexts. <p><i>Optional license: 5 contexts.</i></p>
ASA 5515-X	<p>Base License: 2 contexts.</p> <p><i>Optional license: 5 contexts.</i></p>
ASA 5525-X	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, or 20 contexts.</i></p>
ASA 5545-X	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, or 50 contexts.</i></p>
ASA 5555-X	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, 50, or 100 contexts.</i></p>
ASA 5585-X with SSP-10	<p>Base License: 2 contexts.</p> <p><i>Optional licenses: 5, 10, 20, 50, or 100 contexts.</i></p>

Model	License Requirement
ASA 5585-X with SSP-20, -40, and -60	Base License: 2 contexts. <i>Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.</i>
ASASM	Base License: 2 contexts. <i>Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.</i>

Prerequisites

After you are in multiple context mode, connect to the system or the 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.



Note

Cross context IPv6 routing is not supported.

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 6-9.
- See the “Default MAC Address” section on page 6-12.

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 6-15](#)
- [Enabling or Disabling Multiple Context Mode, page 6-16](#)
- [Configuring a Class for Resource Management, page 6-17](#)
- [Configuring a Security Context, page 6-19](#)
- [Automatically Assigning MAC Addresses to Context Interfaces, page 6-24](#)

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 6-16 . |
| Step 2 | (Optional) Configure classes for resource management. See the “ Configuring a Class for Resource Management ” section on page 6-17 . 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. <ul style="list-style-type: none">• ASA 5500—Chapter 9, “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 6-19 . |
| Step 5 | (Optional) Customize MAC address assignments. See the “ Automatically Assigning MAC Addresses to Context Interfaces ” section on page 6-24 . |

- Step 6** Complete interface configuration in the context. See [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)

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.

This section includes the following topics:

- [Enabling Multiple Context Mode, page 6-16](#)
- [Restoring Single Context Mode, page 6-16](#)

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 [“Backing Up Configurations or Other Files”](#) section on page 42-25.

Detailed Steps

Command	Purpose
<code>mode multiple</code>	Changes to multiple context mode. You are prompted to reboot the ASA.
Example: <code>ciscoasa(config)# mode multiple</code>	

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	<pre>copy disk0:old_running.cfg startup-config</pre> <p>Example:</p> <pre>ciscoasa(config)# copy disk0:old_running.cfg startup-config</pre>	Copies the backup version of your original running configuration to the current startup configuration.
Step 2	<pre>mode single</pre> <p>Example:</p> <pre>ciscoasa(config)# mode single</pre>	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 6-1 lists the resource types and the limits. See also the **show resource types** command.

Table 6-1 Resource Names and Limits

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit ¹	Description
asdm	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.
conns ²	Concurrent or Rate	N/A	Concurrent connections: See the “Supported Feature Licenses Per Model” section on page 4-1 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.

Table 6-1 Resource Names and Limits (continued)

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit ¹	Description
hosts	Concurrent	N/A	N/A	Hosts that can connect through the ASA.
inspects	Rate	N/A	N/A	Application inspections per second.
mac-addresses	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.
vpn burst other	Concurrent	N/A	The Other VPN session amount for your model minus the sum of the sessions assigned to all contexts for vpn other .	The number of site-to-site VPN sessions allowed beyond the amount assigned to a context with vpn other . For example, if your model supports 5000 sessions, and you assign 4000 sessions across all contexts with vpn other , then the remaining 1000 sessions are available for vpn burst other . Unlike vpn other , which guarantees the sessions to the context, vpn burst other can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
vpn other	Concurrent	N/A	See the “Supported Feature Licenses Per Model” section on page 4-1 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.
syslogs	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

	Command	Purpose
Step 1	class <i>name</i> Example: ciscoasa(config)# class gold	Specifies the class name and enters the class configuration mode. The <i>name</i> is a string up to 20 characters long. To set the limits for the default class, enter default for the name.
Step 2	limit-resource [<i>rate</i>] <i>resource_name</i> <i>number</i> [%] Example: ciscoasa(config-class)# limit-resource rate inspects 10	Sets the resource limit for a resource type. See Table 6-1 for a list of resource types. If you specify all , then all resources are configured with the same value. If you also specify a value for a particular resource, the limit overrides the limit set for all . Enter the rate argument to set the rate per second for certain resources. For most resources, specify 0 for the <i>number</i> to set the resource to be unlimited or to be the system limit, if available. For VPN resources, 0 sets the limit to none. For resources that do not have a system limit, you cannot set the percentage (%); you can only set an absolute value.

Examples

For example, to set the default class limit for conns to 10 percent instead of unlimited, and to allow 5 site-to-site VPN tunnels with 2 tunnels allowed for VPN burst, enter the following commands:

```
ciscoasa(config)# class default
ciscoasa(config-class)# limit-resource conns 10%
ciscoasa(config-class)# limit-resource vpn other 5
ciscoasa(config-class)# limit-resource vpn burst other 2
```

All other resources remain at unlimited.

To add a class called gold, enter the following commands:

```
ciscoasa(config)# class gold
ciscoasa(config-class)# limit-resource mac-addresses 10000
ciscoasa(config-class)# limit-resource conns 15%
ciscoasa(config-class)# limit-resource rate conns 1000
ciscoasa(config-class)# limit-resource rate inspects 500
ciscoasa(config-class)# limit-resource hosts 9000
ciscoasa(config-class)# limit-resource asdm 5
ciscoasa(config-class)# limit-resource ssh 5
ciscoasa(config-class)# limit-resource rate syslogs 5000
ciscoasa(config-class)# limit-resource telnet 5
ciscoasa(config-class)# limit-resource xlates 36000
ciscoasa(config-class)# limit-resource routes 5000
ciscoasa(config-class)# limit-resource vpn other 10
ciscoasa(config-class)# limit-resource vpn burst other 5
```

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 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
- If you do not have an admin context (for example, if you clear the configuration) then you must first specify the admin context name by entering the following command:

```
ciscoasa(config)# admin-context name
```

Although this context does not exist yet in your configuration, you can subsequently enter the **context** *name* command to continue the admin context configuration.

Detailed Steps

	Command	Purpose
Step 1	<p>context <i>name</i></p> <p>Example: ciscoasa(config)# context administrator</p>	<p>Adds or modifies a context. The <i>name</i> is a string up to 32 characters long. This name is case sensitive, so you can have two contexts named “customerA” and “CustomerA,” for example. You can use letters, digits, or hyphens, but you cannot start or end the name with a hyphen.</p> <p>“System” or “Null” (in upper or lower case letters) are reserved names, and cannot be used.</p>
Step 2	<p>(Optional)</p> <p>description <i>text</i></p> <p>Example: ciscoasa(config-ctx)# description Administrator Context</p>	<p>Adds a description for this context.</p>

Command	Purpose
<p>Step 3 To allocate an interface:</p> <pre>allocate-interface <i>interface_id</i> [<i>mapped_name</i>] [visible invisible]</pre> <p>To allocate one or more subinterfaces:</p> <pre>allocate-interface <i>interface_id.subinterface</i>[-<i>interface_id.subinterface</i>] [<i>mapped_name</i>[-<i>mapped_name</i>]] [visible invisible]</pre> <p>Example:</p> <pre>ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.100 int1 ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.200 int2 ciscoasa(config-ctx)# allocate-interface gigabitethernet0/2.300-gigabitethernet0/2. 305 int3-int8</pre>	<p>Specifies the interfaces you can use in the context. Do not include a space between the interface type and the port number.</p> <p>Enter these commands multiple times to specify different ranges. If you remove an allocation with the no form of this command, then any context commands that include this interface are removed from the running configuration.</p> <p>Transparent firewall mode allows a limited number of interfaces to pass through traffic; however, you can use a dedicated management interface, Management <i>slot/port</i> (physical, subinterface, redundant, or EtherChannel), as an additional interface for management traffic. A separate management interface is not available for the ASASM.</p> <p>You can assign the same interfaces to multiple contexts in routed mode, if desired. Transparent mode does not allow shared interfaces.</p> <p>The <i>mapped_name</i> is an alphanumeric alias for the interface that can be used within the context instead of the interface ID. If you do not specify a mapped name, the interface ID is used within the context. For security purposes, you might not want the context administrator to know which interfaces the context is using. A mapped name must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, or an underscore. For example, you can use the following names:</p> <pre>int0, inta, int_0</pre> <p>If you specify a range of subinterfaces, you can specify a matching range of mapped names. Follow these guidelines for ranges:</p> <ul style="list-style-type: none"> The mapped name must consist of an alphabetic portion followed by a numeric portion. The alphabetic portion of the mapped name must match for both ends of the range. For example, enter the following range: <pre>int0-int10</pre> <p>If you enter <code>gig0/1.1-gig0/1.5 happy1-sad5</code>, for example, the command fails.</p> The numeric portion of the mapped name must include the same quantity of numbers as the subinterface range. For example, both ranges include 100 interfaces: <pre>gigabitethernet0/0.100-gigabitethernet0/0.199 int1-int100</pre> <p>If you enter <code>gig0/0.100-gig0/0.199 int1-int15</code>, for example, the command fails.</p> <p>Specify visible to see the real interface ID in the show interface command if you set a mapped name. The default invisible keyword shows only the mapped name.</p>

Command	Purpose
<p>Step 4</p> <p>config-url <i>url</i></p> <p>Example: <pre>ciscoasa(config-ctx)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/te st.cfg</pre></p>	<p>Identifies the URL from which the system downloads the context configuration. When you add a context URL, the system immediately loads the context so that it is running, if the configuration is available.</p> <p>Note Enter the allocate-interface command(s) before you enter the config-url command. If you enter the config-url command first, the ASA loads the context configuration immediately. If the context contains any commands that refer to (not yet configured) interfaces, those commands fail.</p> <p>The filename does not require a file extension, although we recommend using “.cfg”. The server must be accessible from the admin context. If the configuration file is not available, you see the following message:</p> <pre>WARNING: Could not fetch the URL url INFO: Creating context with default config</pre> <p>For non-HTTP(S) URL locations, after you specify the URL, you can then change to the context, configure it at the CLI, and enter the write memory command to write the file to the URL location. (HTTP(S) is read only).</p> <p>Note The admin context file must be stored on the internal flash memory.</p> <p>Available URL types include: disknumber (for flash memory), ftp, http, https, or fttp.</p> <p>To change the URL, reenter the config-url command with a new URL. See the “Changing the Security Context URL” section on page 6-26 for more information about changing the URL.</p>
<p>Step 5</p> <p>(Optional)</p> <p>member <i>class_name</i></p> <p>Example: <pre>ciscoasa(config-ctx)# member gold</pre></p>	<p>Assigns the context to a resource class. If you do not specify a class, the context belongs to the default class. You can only assign a context to one resource class.</p>
<p>Step 6</p> <p>(Optional)</p> <p>allocate-ips <i>sensor_name</i> [<i>mapped_name</i>] [default]</p> <p>Example: <pre>ciscoasa(config-ctx)# allocate-ips sensor1 highsec</pre></p>	<p>Assigns an IPS virtual sensor to this context if you have the IPS module installed.</p> <p>See the Assigning Virtual Sensors to a Security Context (ASA 5510 and Higher) section in the firewall configuration guide for detailed information about virtual sensors.</p>

	Command	Purpose
Step 7	(Optional) join-failover-group {1 2} Example: ciscoasa(config-ctx)# join-failover-group 2	Assigns a context to a failover group in Active/Active failover. By default, contexts are in group 1. The admin context must always be in group 1. See the “Configuring Optional Failover Parameters” section on page 7-35 for detailed information about failover groups.
Step 8	(Optional) scansafe [license key] Example: ciscoasa(config-ctx)# scansafe	Enables Cloud Web Security for this context. If you do not specify a license , the context uses the license configured in the system configuration. The ASA sends the authentication key to the Cloud Web Security proxy servers to indicate from which organization the request comes. The authentication key is a 16-byte hexadecimal number. See the Configuring the ASA for Cisco Cloud Web Security section in the firewall configuration guide for detailed information about ScanSafe.

Examples

The following example sets the admin context to be “administrator,” creates a context called “administrator” on the internal flash memory, and then adds two contexts from an FTP server:

```

ciscoasa(config)# admin-context administrator
ciscoasa(config)# context administrator
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.1
ciscoasa(config-ctx)# config-url disk0:/admin.cfg

ciscoasa(config-ctx)# context test
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.100 int1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.102 int2
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.110-gigabitethernet0/0.115
int3-int8
ciscoasa(config-ctx)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/test.cfg
ciscoasa(config-ctx)# member gold

ciscoasa(config-ctx)# context sample
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.200 int1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.212 int2
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.230-gigabitethernet0/1.235
int3-int8
ciscoasa(config-ctx)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/sample.cfg
ciscoasa(config-ctx)# member silver

```

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 6-11](#) for more information, especially if you are upgrading from an earlier ASA version. See also the [“Viewing Assigned MAC Addresses” section on page 6-36](#).

Guidelines

- When you configure a **nameif** command 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 11-10](#) to manually set the MAC address.

Detailed Steps

Command	Purpose
mac-address auto [<i>prefix prefix</i>] Example: <pre>ciscoasa(config)# mac-address auto prefix 19</pre>	<p>Automatically assigns private MAC addresses to each context interface.</p> <p>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.</p> <p>If you manually enter a prefix, then the <i>prefix</i> is 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 6-12 section for more information about how the prefix is used.</p>

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, or that is used in the **copy** or **write** commands, 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. For example, you cannot view all running configurations (system plus all contexts) by entering the **show running-config** command. Only the current configuration displays.

Detailed Steps

Command	Purpose
<code>changeto context name</code>	Changes to a context. The prompt changes to the following: ciscoasa/name#
<code>changeto system</code>	Changes to the system execution space. The prompt changes to the following: ciscoasa#

Managing Security Contexts

This section describes how to manage security contexts and includes the following topics:

- [Removing a Security Context, page 6-25](#)
- [Changing the Admin Context, page 6-26](#)
- [Changing the Security Context URL, page 6-26](#)
- [Reloading a Security Context, page 6-27](#)

Removing a Security Context

You cannot remove the current admin context, unless you remove all contexts using the **clear context** command.



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. You might see an error message indicating that the number of interfaces on the active and standby units are not consistent; this error is temporary and can be ignored.

Prerequisites

Perform this procedure in the system execution space.

Detailed Steps

Command	Purpose
<code>no context name</code>	Removes a single context. All context commands are also removed. The context configuration file is not removed from the config URL location.
<code>clear context</code>	Removes all contexts (including the admin context). The context configuration files are not removed from the config URL locations.

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.

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

Command	Purpose
admin-context <i>context_name</i> Example: <pre>ciscoasa(config)# admin-context administrator</pre>	Sets the admin context. Any remote management sessions, such as Telnet, SSH, or HTTPS, 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

	Command	Purpose
Step 1	<p>(Optional, if you do not want to perform a merge)</p> <pre>changeto context name clear configure all</pre> <p>Example: ciscoasa(config)# changeto context ctx1 ciscoasa/ctx1(config)# clear configure all</p>	Changes to the context and clears its configuration. If you want to perform a merge, skip to Step 2.
Step 2	<pre>changeto system</pre> <p>Example: ciscoasa/ctx1(config)# changeto system ciscoasa(config)#</p>	Changes to the system execution space.
Step 3	<pre>context name</pre> <p>Example: ciscoasa(config)# context ctx1</p>	Enters the context configuration mode for the context you want to change.
Step 4	<pre>config-url new_url</pre> <p>Example: ciscoasa(config)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/c tx1.cfg</p>	Enters the new URL. 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 6-28](#)

- [Reloading by Removing and Re-adding the Context, page 6-28](#)

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

	Command	Purpose
Step 1	<p>changeto context <i>name</i></p> <p>Example: ciscoasa(config)# changeto context ctx1 ciscoasa/ctx1(comfig)#</p>	Changes to the context that you want to reload.
Step 2	<p>clear configure all</p> <p>Example: ciscoasa/ctx1(config)# clear configure all</p>	Clears the running configuration. This command clears all connections.
Step 3	<p>copy startup-config running-config</p> <p>Example: ciscoasa/ctx1(config)# copy startup-config running-config</p>	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 6-25.](#)
2. [“Configuring a Security Context” section on page 6-19](#)

Monitoring Security Contexts

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

- [Viewing Context Information, page 6-29](#)
- [Viewing Resource Allocation, page 6-30](#)
- [Viewing Resource Usage, page 6-33](#)
- [Monitoring SYN Attacks in Contexts, page 6-34](#)
- [Viewing Assigned MAC Addresses, page 6-36](#)

Viewing Context Information

From the system execution space, you can view a list of contexts including the name, allocated interfaces, and configuration file URL.

From the system execution space, view all contexts by entering the following command:

Command	Purpose
<code>show context [name detail count]</code>	Shows all contexts. If you want to show information for a particular context, specify the <i>name</i> . The detail option shows additional information. See the following sample outputs below for more information. The count option shows the total number of contexts.

The following is sample output from the **show context** command. The following sample output shows three contexts:

```
ciscoasa# show context

Context Name      Interfaces                                URL
*admin            GigabitEthernet0/1.100                   disk0:/admin.cfg
                  GigabitEthernet0/1.101
contexta         GigabitEthernet0/1.200                   disk0:/contexta.cfg
                  GigabitEthernet0/1.201
contextb         GigabitEthernet0/1.300                   disk0:/contextb.cfg
                  GigabitEthernet0/1.301
Total active Security Contexts: 3
```

Table 6-2 shows each field description.

Table 6-2 *show context Fields*

Field	Description
Context Name	Lists all context names. The context name with the asterisk (*) is the admin context.
Interfaces	The interfaces assigned to the context.
URL	The URL from which the ASA loads the context configuration.

The following is sample output from the **show context detail** command:

```
ciscoasa# show context detail

Context "admin", has been created, but initial ACL rules not complete
Config URL: disk0:/admin.cfg
Real Interfaces: Management0/0
Mapped Interfaces: Management0/0
Flags: 0x00000013, ID: 1

Context "ctx", has been created, but initial ACL rules not complete
Config URL: ctx.cfg
Real Interfaces: GigabitEthernet0/0.10, GigabitEthernet0/1.20,
                GigabitEthernet0/2.30
Mapped Interfaces: int1, int2, int3
Flags: 0x00000011, ID: 2
```

```

Context "system", is a system resource
  Config URL: startup-config
  Real Interfaces:
  Mapped Interfaces: Control0/0, GigabitEthernet0/0,
    GigabitEthernet0/0.10, GigabitEthernet0/1, GigabitEthernet0/1.10,
    GigabitEthernet0/1.20, GigabitEthernet0/2, GigabitEthernet0/2.30,
    GigabitEthernet0/3, Management0/0, Management0/0.1
  Flags: 0x00000019, ID: 257

Context "null", is a system resource
  Config URL: ... null ...
  Real Interfaces:
  Mapped Interfaces:
  Flags: 0x00000009, ID: 258

```

See the command reference for more information about the **detail** output.

The following is sample output from the **show context count** command:

```

ciscoasa# show context count
Total active contexts: 2

```

Viewing Resource Allocation

From the system execution space, you can view the allocation for each resource across all classes and class members.

To view the resource allocation, enter the following command:

Command	Purpose
<code>show resource allocation [detail]</code>	Shows the resource allocation. This command shows the resource allocation, but does not show the actual resources being used. See the “Viewing Resource Usage” section on page 6-33 for more information about actual resource usage. The detail argument shows additional information. See the following sample outputs for more information.

The following sample output shows the total allocation of each resource as an absolute value and as a percentage of the available system resources:

```

ciscoasa# show resource allocation
Resource          Total          % of Avail
Conns [rate]      35000         N/A
Inspects [rate]   35000         N/A
Syslogs [rate]    10500         N/A
Conns              305000        30.50%
Hosts              78842         N/A
SSH                35            35.00%
Routes             5000          N/A
Telnet             35            35.00%
Xlates             91749         N/A
Other VPN Sessions 20            2.66%
Other VPN Burst    20            2.66%
All                unlimited

```


Table 6-3 shows each field description.

Table 6-3 *show resource allocation Fields*

Field	Description
Resource	The name of the resource that you can limit.
Total	The total amount of the resource that is allocated across all contexts. The amount is an absolute number of concurrent instances or instances per second. If you specified a percentage in the class definition, the ASA converts the percentage to an absolute number for this display.
% of Avail	The percentage of the total system resources that is allocated across all contexts, if the resource has a hard system limit. If a resource does not have a system limit, this column shows N/A.

The following is sample output from the **show resource allocation detail** command:

```
ciscoasa# show resource allocation detail
Resource Origin:
  A Value was derived from the resource 'all'
  C Value set in the definition of this class
  D Value set in default class
Resource      Class      Mmbrs  Origin  Limit      Total      Total %
Conns [rate]  default    all    CA      unlimited
              gold       1      C       34000     34000     N/A
              silver    1      CA      17000     17000     N/A
              bronze   0      CA       8500
              All Contexts: 3              51000     N/A

Inspects [rate] default    all    CA      unlimited
              gold       1      DA      unlimited
              silver    1      CA      10000     10000     N/A
              bronze   0      CA       5000
              All Contexts: 3              10000     N/A

Syslogs [rate] default    all    CA      unlimited
              gold       1      C       6000      6000      N/A
              silver    1      CA      3000      3000      N/A
              bronze   0      CA      1500
              All Contexts: 3              9000      N/A

Conns         default    all    CA      unlimited
              gold       1      C       200000    200000    20.00%
              silver    1      CA      100000    100000    10.00%
              bronze   0      CA      50000
              All Contexts: 3              300000    30.00%

Hosts        default    all    CA      unlimited
              gold       1      DA      unlimited
              silver    1      CA      26214     26214     N/A
              bronze   0      CA      13107
              All Contexts: 3              26214     N/A

SSH          default    all    C       5
              gold       1      D       5          5          5.00%
              silver    1      CA      10         10         10.00%
              bronze   0      CA       5
              All Contexts: 3              20         20.00%

Telnet       default    all    C       5
```

	gold	1	D	5	5	5.00%
	silver	1	CA	10	10	10.00%
	bronze	0	CA	5		
	All Contexts:	3			20	20.00%
Routes	default	all	C	unlimited		N/A
	gold	1	D	unlimited	5	N/A
	silver	1	CA	10	10	N/A
	bronze	0	CA	5		N/A
	All Contexts:	3			20	N/A
Xlates	default	all	CA	unlimited		
	gold	1	DA	unlimited		
	silver	1	CA	23040	23040	N/A
	bronze	0	CA	11520		
	All Contexts:	3			23040	N/A
mac-addresses	default	all	C	65535		
	gold	1	D	65535	65535	100.00%
	silver	1	CA	6553	6553	9.99%
	bronze	0	CA	3276		
	All Contexts:	3			137623	209.99%

Table 6-4 shows each field description.

Table 6-4 show resource allocation detail Fields

Field	Description
Resource	The name of the resource that you can limit.
Class	The name of each class, including the default class. The All contexts field shows the total values across all classes.
Mmbrs	The number of contexts assigned to each class.
Origin	The origin of the resource limit, as follows: <ul style="list-style-type: none"> A—You set this limit with the all option, instead of as an individual resource. C—This limit is derived from the member class. D—This limit was not defined in the member class, but was derived from the default class. For a context assigned to the default class, the value will be “C” instead of “D.” The ASA can combine “A” with “C” or “D.”
Limit	The limit of the resource per context, as an absolute number. If you specified a percentage in the class definition, the ASA converts the percentage to an absolute number for this display.
Total	The total amount of the resource that is allocated across all contexts in the class. The amount is an absolute number of concurrent instances or instances per second. If the resource is unlimited, this display is blank.
% of Avail	The percentage of the total system resources that is allocated across all contexts in the class. If the resource is unlimited, this display is blank. If the resource does not have a system limit, then this column shows N/A.

Viewing Resource Usage

From the system execution space, you can view the resource usage for each context and display the system resource usage.

Command	Purpose
<pre>show resource usage [context context_name top n all summary system] [resource {resource_name all} detail] [counter counter_name [count_threshold]]</pre>	<p>By default, all context usage is displayed; each context is listed separately.</p> <p>Enter the top n keyword to show the contexts that are the top <i>n</i> users of the specified resource. You must specify a single resource type, and not resource all, with this option.</p> <p>The summary option shows all context usage combined.</p> <p>The system option shows all context usage combined, but shows the system limits for resources instead of the combined context limits.</p> <p>For the resource resource_name, see Table 6-1 for available resource names. See also the show resource type command. Specify all (the default) for all types.</p> <p>The detail option shows the resource usage of all resources, including those you cannot manage. For example, you can view the number of TCP intercepts.</p> <p>The counter counter_name is one of the following keywords:</p> <ul style="list-style-type: none"> • current—Shows the active concurrent instances or the current rate of the resource. • denied—Shows the number of instances that were denied because they exceeded the resource limit shown in the Limit column. • peak—Shows the peak concurrent instances, or the peak rate of the resource since the statistics were last cleared, either using the clear resource usage command or because the device rebooted. • all—(Default) Shows all statistics. <p>The <i>count_threshold</i> sets the number above which resources are shown. The default is 1. If the usage of the resource is below the number you set, then the resource is not shown. If you specify all for the counter name, then the <i>count_threshold</i> applies to the current usage.</p> <p>Note To show all resources, set the <i>count_threshold</i> to 0.</p>

The following is sample output from the **show resource usage context** command, which shows the resource usage for the admin context:

```
ciscoasa# show resource usage context admin
```

Resource	Current	Peak	Limit	Denied	Context
Telnet	1	1	5	0	admin
Conns	44	55	N/A	0	admin
Hosts	45	56	N/A	0	admin

The following is sample output from the **show resource usage summary** command, which shows the resource usage for all contexts and all resources. This sample shows the limits for six contexts.

```
ciscoasa# show resource usage summary
```

Resource	Current	Peak	Limit	Denied	Context
Syslogs [rate]	1743	2132	N/A	0	Summary
Conns	584	763	280000 (S)	0	Summary
Xlates	8526	8966	N/A	0	Summary
Hosts	254	254	N/A	0	Summary
Conns [rate]	270	535	N/A	1704	Summary
Inspects [rate]	270	535	N/A	0	Summary
Other VPN Sessions	0	10	10	740	Summary
Other VPN Burst	0	10	10	730	Summary

S = System: Combined context limits exceed the system limit; the system limit is shown.

The following is sample output from the **show resource usage summary** command, which shows the limits for 25 contexts. Because the context limit for Telnet and SSH connections is 5 per context, then the combined limit is 125. The system limit is only 100, so the system limit is shown.

```
ciscoasa# show resource usage summary
```

Resource	Current	Peak	Limit	Denied	Context
Telnet	1	1	100 [S]	0	Summary
SSH	2	2	100 [S]	0	Summary
Conns	56	90	130000 (S)	0	Summary
Hosts	89	102	N/A	0	Summary

S = System: Combined context limits exceed the system limit; the system limit is shown.

The following is sample output from the **show resource usage system** command, which shows the resource usage for all contexts, but it shows the system limit instead of the combined context limits. The **counter all 0** option is used to show resources that are not currently in use. The Denied statistics indicate how many times the resource was denied due to the system limit, if available.

```
ciscoasa# show resource usage system counter all 0
```

Resource	Current	Peak	Limit	Denied	Context
Telnet	0	0	100	0	System
SSH	0	0	100	0	System
ASDM	0	0	32	0	System
Routes	0	0	N/A	0	System
IPSec	0	0	5	0	System
Syslogs [rate]	1	18	N/A	0	System
Conns	0	1	280000	0	System
Xlates	0	0	N/A	0	System
Hosts	0	2	N/A	0	System
Conns [rate]	1	1	N/A	0	System
Inspects [rate]	0	0	N/A	0	System
Other VPN Sessions	0	10	750	740	System
Other VPN Burst	0	10	750	730	System

Monitoring SYN Attacks in Contexts

The ASA prevents SYN attacks using TCP Intercept. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the ASA acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the ASA receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

Monitor SYN attacks using the following commands:

Command	Purpose
<code>show perfmon</code>	Monitors the rate of attacks for individual contexts.
<code>show resource usage detail</code>	Monitors the amount of resources being used by TCP intercept for individual contexts.
<code>show resource usage summary detail</code>	Monitors the resources being used by TCP intercept for the entire system.

The following is sample output from the `show perfmon` command that shows the rate of TCP intercepts for a context called admin.

```
ciscoasa/admin# show perfmon

Context: admin
PERFMON STATS:   Current      Average
Xlates           0/s          0/s
Connections      0/s          0/s
TCP Conns        0/s          0/s
UDP Conns        0/s          0/s
URL Access       0/s          0/s
URL Server Req   0/s          0/s
WebSns Req       0/s          0/s
TCP Fixup        0/s          0/s
HTTP Fixup       0/s          0/s
FTP Fixup        0/s          0/s
AAA Authen       0/s          0/s
AAA Author       0/s          0/s
AAA Account      0/s          0/s
TCP Intercept    322779/s     322779/s
```

The following is sample output from the `show resource usage detail` command that shows the amount of resources being used by TCP Intercept for individual contexts. (Sample text in **bold** shows the TCP intercept information.)

```
ciscoasa(config)# show resource usage detail

Resource          Current      Peak      Limit      Denied Context
memory            843732      847288   unlimited  0 admin
chunk:channels    14          15       unlimited  0 admin
chunk:fixup       15          15       unlimited  0 admin
chunk:hole        1           1        unlimited  0 admin
chunk:ip-users    10          10       unlimited  0 admin
chunk:list-elem   21          21       unlimited  0 admin
chunk:list-hdr    3           4        unlimited  0 admin
chunk:route       2           2        unlimited  0 admin
chunk:static      1           1        unlimited  0 admin
tcp-intercepts   328787     803610  unlimited  0 admin
np-statics        3           3        unlimited  0 admin
statics           1           1        unlimited  0 admin
ace-rules         1           1        unlimited  0 admin
console-access-rul 2           2        unlimited  0 admin
fixup-rules       14          15       unlimited  0 admin
memory            959872      960000   unlimited  0 c1
chunk:channels    15          16       unlimited  0 c1
chunk:dbgtrace    1           1        unlimited  0 c1
chunk:fixup       15          15       unlimited  0 c1
chunk:global      1           1        unlimited  0 c1
chunk:hole        2           2        unlimited  0 c1
chunk:ip-users    10          10       unlimited  0 c1
chunk:udp-ctrl-blk 1           1        unlimited  0 c1
```

```

chunk:list-elem          24          24 unlimited          0 c1
chunk:list-hdr          5           6 unlimited          0 c1
chunk:nat                1           1 unlimited          0 c1
chunk:route             2           2 unlimited          0 c1
chunk:static            1           1 unlimited          0 c1
tcp-intercept-rate    16056      16254 unlimited          0 c1
globals                 1           1 unlimited          0 c1
np-statics              3           3 unlimited          0 c1
statics                 1           1 unlimited          0 c1
nats                    1           1 unlimited          0 c1
ace-rules               2           2 unlimited          0 c1
console-access-rul     2           2 unlimited          0 c1
fixup-rules            14          15 unlimited          0 c1
memory                  232695716  232020648 unlimited          0 system
chunk:channels          17          20 unlimited          0 system
chunk:dbgtrace          3           3 unlimited          0 system
chunk:fixup            15          15 unlimited          0 system
chunk:ip-users          4           4 unlimited          0 system
chunk:list-elem        1014        1014 unlimited          0 system
chunk:list-hdr          1           1 unlimited          0 system
chunk:route             1           1 unlimited          0 system
block:16384             510         885 unlimited          0 system
block:2048              32          34 unlimited          0 system

```

The following sample output shows the resources being used by TCP intercept for the entire system. (Sample text in **bold** shows the TCP intercept information.)

```

ciscoasa(config)# show resource usage summary detail
Resource           Current      Peak      Limit      Denied Context
memory             238421312  238434336 unlimited  0 Summary
chunk:channels     46         48 unlimited  0 Summary
chunk:dbgtrace     4           4 unlimited  0 Summary
chunk:fixup        45         45 unlimited  0 Summary
chunk:global       1           1 unlimited  0 Summary
chunk:hole         3           3 unlimited  0 Summary
chunk:ip-users     24         24 unlimited  0 Summary
chunk:udp-ctrl-blk 1           1 unlimited  0 Summary
chunk:list-elem    1059       1059 unlimited  0 Summary
chunk:list-hdr     10         11 unlimited  0 Summary
chunk:nat          1           1 unlimited  0 Summary
chunk:route        5           5 unlimited  0 Summary
chunk:static       2           2 unlimited  0 Summary
block:16384        510        885 unlimited  0 Summary
block:2048         32         35 unlimited  0 Summary
tcp-intercept-rate 341306    811579 unlimited 0 Summary
globals            1           1 unlimited  0 Summary
np-statics         6           6 unlimited  0 Summary
statics            2           2          N/A        0 Summary
nats               1           1          N/A        0 Summary
ace-rules          3           3          N/A        0 Summary
console-access-rul 4           4          N/A        0 Summary
fixup-rules        43         44          N/A        0 Summary

```

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 6-37](#)
- [Viewing MAC Addresses Within a Context, page 6-38](#)

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

Command	Purpose
<code>show running-config all context [name]</code>	Shows the assigned MAC addresses from the system execution space. The all option is required to view the assigned MAC addresses. Although the mac-address auto command is user-configurable in global configuration mode only, the command appears as a read-only entry in context configuration mode along with the assigned MAC address. Only allocated interfaces that are configured with a nameif command within the context have a MAC address assigned.

Examples

The following output from the **show running-config all context admin** command shows the primary and standby MAC address assigned to the Management0/0 interface:

```
ciscoasa# show running-config all context admin

context admin
  allocate-interface Management0/0
  mac-address auto Management0/0 a24d.0000.1440 a24d.0000.1441
  config-url disk0:/admin.cfg
```

The following output from the **show running-config all context** command shows all the MAC addresses (primary and standby) for all context interfaces. Note that because the GigabitEthernet0/0 and GigabitEthernet0/1 main interfaces are not configured with a **nameif** command inside the contexts, no MAC addresses have been generated for them.

```
ciscoasa# show running-config all context

admin-context admin
context admin
  allocate-interface Management0/0
  mac-address auto Management0/0 a2d2.0400.125a a2d2.0400.125b
  config-url disk0:/admin.cfg
!

context CTX1
  allocate-interface GigabitEthernet0/0
  allocate-interface GigabitEthernet0/0.1-GigabitEthernet0/0.5
  mac-address auto GigabitEthernet0/0.1 a2d2.0400.11bc a2d2.0400.11bd
  mac-address auto GigabitEthernet0/0.2 a2d2.0400.11c0 a2d2.0400.11c1
  mac-address auto GigabitEthernet0/0.3 a2d2.0400.11c4 a2d2.0400.11c5
  mac-address auto GigabitEthernet0/0.4 a2d2.0400.11c8 a2d2.0400.11c9
```

```

mac-address auto GigabitEthernet0/0.5 a2d2.0400.11cc a2d2.0400.11cd
allocate-interface GigabitEthernet0/1
allocate-interface GigabitEthernet0/1.1-GigabitEthernet0/1.3
mac-address auto GigabitEthernet0/1.1 a2d2.0400.120c a2d2.0400.120d
mac-address auto GigabitEthernet0/1.2 a2d2.0400.1210 a2d2.0400.1211
mac-address auto GigabitEthernet0/1.3 a2d2.0400.1214 a2d2.0400.1215
config-url disk0:/CTX1.cfg
!

context CTX2
allocate-interface GigabitEthernet0/0
allocate-interface GigabitEthernet0/0.1-GigabitEthernet0/0.5
mac-address auto GigabitEthernet0/0.1 a2d2.0400.11ba a2d2.0400.11bb
mac-address auto GigabitEthernet0/0.2 a2d2.0400.11be a2d2.0400.11bf
mac-address auto GigabitEthernet0/0.3 a2d2.0400.11c2 a2d2.0400.11c3
mac-address auto GigabitEthernet0/0.4 a2d2.0400.11c6 a2d2.0400.11c7
mac-address auto GigabitEthernet0/0.5 a2d2.0400.11ca a2d2.0400.11cb
allocate-interface GigabitEthernet0/1
allocate-interface GigabitEthernet0/1.1-GigabitEthernet0/1.3
mac-address auto GigabitEthernet0/1.1 a2d2.0400.120a a2d2.0400.120b
mac-address auto GigabitEthernet0/1.2 a2d2.0400.120e a2d2.0400.120f
mac-address auto GigabitEthernet0/1.3 a2d2.0400.1212 a2d2.0400.1213
config-url disk0:/CTX2.cfg
!
```

Viewing MAC Addresses Within a Context

This section describes how to view MAC addresses within a context.

Detailed Steps

Command	Purpose
<code>show interface include (Interface) (MAC)</code>	Shows the MAC address in use by each interface within the context.

Examples

For example:

```

ciscoasa/context# show interface | include (Interface)|(MAC)

Interface GigabitEthernet1/1.1 "g1/1.1", is down, line protocol is down
MAC address a201.0101.0600, MTU 1500
Interface GigabitEthernet1/1.2 "g1/1.2", is down, line protocol is down
MAC address a201.0102.0600, MTU 1500
Interface GigabitEthernet1/1.3 "g1/1.3", is down, line protocol is down
MAC address a201.0103.0600, MTU 1500
...
```



Note

The `show interface` command 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.

Configuration Examples for Multiple Context Mode

The following example:

- Automatically sets the MAC addresses in contexts with a custom prefix.
- Sets the default class limit for conns to 10 percent instead of unlimited, and sets the VPN other sessions to 10, with a burst of 5.
- Creates a gold resource class.
- Sets the admin context to be “administrator.”
- Creates a context called “administrator” on the internal flash memory to be part of the default resource class.
- Adds two contexts from an FTP server as part of the gold resource class.

```

ciscoasa(config)# mac-address auto prefix 19

ciscoasa(config)# class default
ciscoasa(config-class)# limit-resource conns 10%
ciscoasa(config-class)# limit-resource vpn other 10
ciscoasa(config-class)# limit-resource vpn burst other 5

ciscoasa(config)# class gold
ciscoasa(config-class)# limit-resource mac-addresses 10000
ciscoasa(config-class)# limit-resource conns 15%
ciscoasa(config-class)# limit-resource rate conns 1000
ciscoasa(config-class)# limit-resource rate inspects 500
ciscoasa(config-class)# limit-resource hosts 9000
ciscoasa(config-class)# limit-resource asdm 5
ciscoasa(config-class)# limit-resource ssh 5
ciscoasa(config-class)# limit-resource rate syslogs 5000
ciscoasa(config-class)# limit-resource telnet 5
ciscoasa(config-class)# limit-resource xlates 36000
ciscoasa(config-class)# limit-resource routes 700
ciscoasa(config-class)# limit-resource vpn other 100
ciscoasa(config-class)# limit-resource vpn burst other 50

ciscoasa(config)# admin-context administrator
ciscoasa(config)# context administrator
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.1
ciscoasa(config-ctx)# config-url disk0:/admin.cfg

ciscoasa(config-ctx)# context test
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.100 int1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.102 int2
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/0.110-gigabitethernet0/0.115
int3-int8
ciscoasa(config-ctx)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/test.cfg
ciscoasa(config-ctx)# member gold

ciscoasa(config-ctx)# context sample
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.200 int1
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.212 int2
ciscoasa(config-ctx)# allocate-interface gigabitethernet0/1.230-gigabitethernet0/1.235
int3-int8
ciscoasa(config-ctx)# config-url ftp://user1:passw0rd@10.1.1.1/configlets/sample.cfg
ciscoasa(config-ctx)# member gold

```

Feature History for Multiple Context Mode

Table 6-5 lists each feature change and the platform release in which it was implemented.

Table 6-5 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 commands: context , mode , and class .
Automatic MAC address assignment	7.2(1)	Automatic assignment of MAC address to context interfaces was introduced. We introduced the following command: mac-address auto .
Resource management	7.2(1)	Resource management was introduced. We introduced the following commands: class , limit-resource , and member .
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 introduced the following command: allocate-ips .
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 command: mac-address auto prefix .
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 command: mac-address auto .

Table 6-5 Feature History for Multiple Context Mode (continued)

Feature Name	Platform Releases	Feature Information
Automatic generation of a MAC address prefix	8.6(1)	<p>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 reenables MAC address generation. The prefix method of generation provides many benefits, including a better guarantee of unique MAC addresses on a segment. You can view the auto-generated prefix by entering the show running-config mac-address command. 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.</p> <p>Note To maintain hitless upgrade for failover pairs, the ASA does <i>not</i> 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, reenables MAC address generation to use the default prefix.</p> <p>We modified the following command: mac-address auto.</p>
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)	<p>A new resource type, routes, was created to set the maximum number of routing table entries in each context.</p> <p>We modified the following commands: limit-resource, show resource types, show resource usage, show resource allocation.</p>

Table 6-5 Feature History for Multiple Context Mode (continued)

Feature Name	Platform Releases	Feature Information
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)	<p>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.</p> <p>We modified the following commands: limit-resource, show resource types, show resource usage, show resource allocation.</p>



Configuring Failover

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

- [Introduction to Failover, page 7-1](#)
- [Licensing Requirements Failover, page 7-24](#)
- [Prerequisites for Failover, page 7-24](#)
- [Guidelines and Limitations, page 7-25](#)
- [Default Settings, page 7-25](#)
- [Configuring Active/Standby Failover, page 7-26](#)
- [Configuring Active/Active Failover, page 7-30](#)
- [Configuring Optional Failover Parameters, page 7-35](#)
- [Managing Failover, page 7-42](#)
- [Monitoring Failover, page 7-48](#)
- [Feature History for Failover, page 7-49](#)

Introduction to Failover

- [Failover Overview, page 7-2](#)
- [Failover System Requirements, page 7-2](#)
- [Failover and Stateful Failover Links, page 7-3](#)
- [MAC Addresses and IP Addresses, page 7-7](#)
- [Intra- and Inter-Chassis Module Placement for the ASA Services Module, page 7-8](#)
- [Stateless and Stateful Failover, page 7-12](#)
- [Transparent Firewall Mode Requirements, page 7-14](#)
- [Failover Health Monitoring, page 7-16](#)
- [Failover Times, page 7-18](#)
- [Configuration Synchronization, page 7-18](#)
- [Information About Active/Standby Failover, page 7-20](#)
- [Information About Active/Active Failover, page 7-21](#)

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 7-2](#)
- [Software Requirements, page 7-2](#)
- [License Requirements, page 7-3](#)

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 42-5](#) 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 4-30](#) 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 7-3](#)
- [Stateful Failover Link, page 7-4](#)
- [Avoiding Interrupted Failover and Data Links, page 7-5](#)



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 7-3](#)
- [Interface for the Failover Link, page 7-4](#)
- [Connecting the Failover Link, page 7-4](#)

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 7-4](#)
- [Shared with the Failover Link, page 7-5](#)
- [Shared with a Regular Data Interface \(Not Recommended\), page 7-5](#)

**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 than 10 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 7-1](#) and [Figure 7-2](#) are NOT recommended.

Figure 7-1 Connecting with a Single Switch—Not Recommended



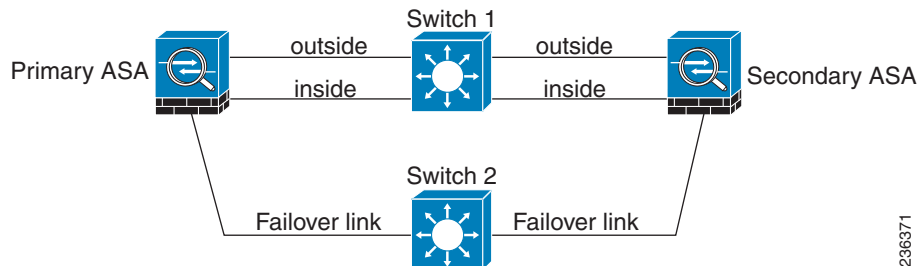



Figure 7-2 Connecting with a Double Switch—Not Recommended

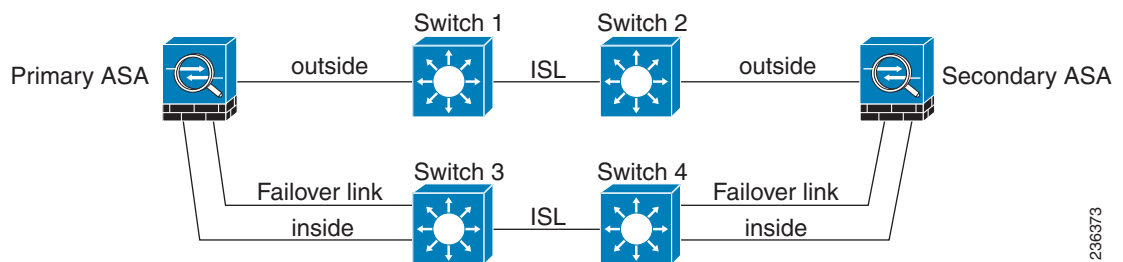


Scenario 2—Recommended

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 7-3](#) and [Figure 7-4](#).

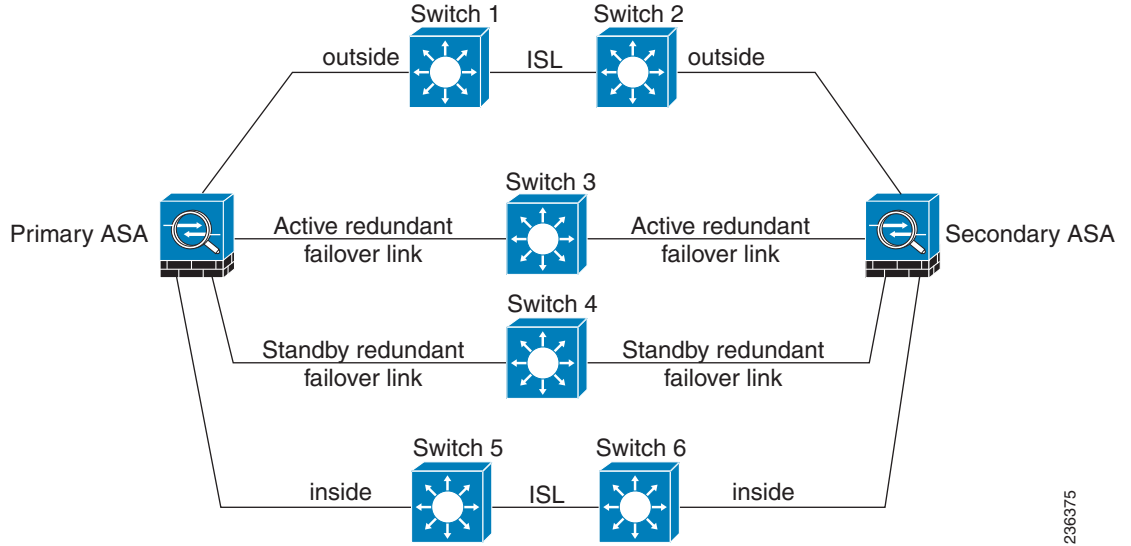
Figure 7-3 *Connecting with a Different Switch***Figure 7-4** *Connecting with a Cable***Scenario 3—Recommended**

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 7-5](#).

Figure 7-5 *Connecting with a Secure Switch***Scenario 4—Recommended**

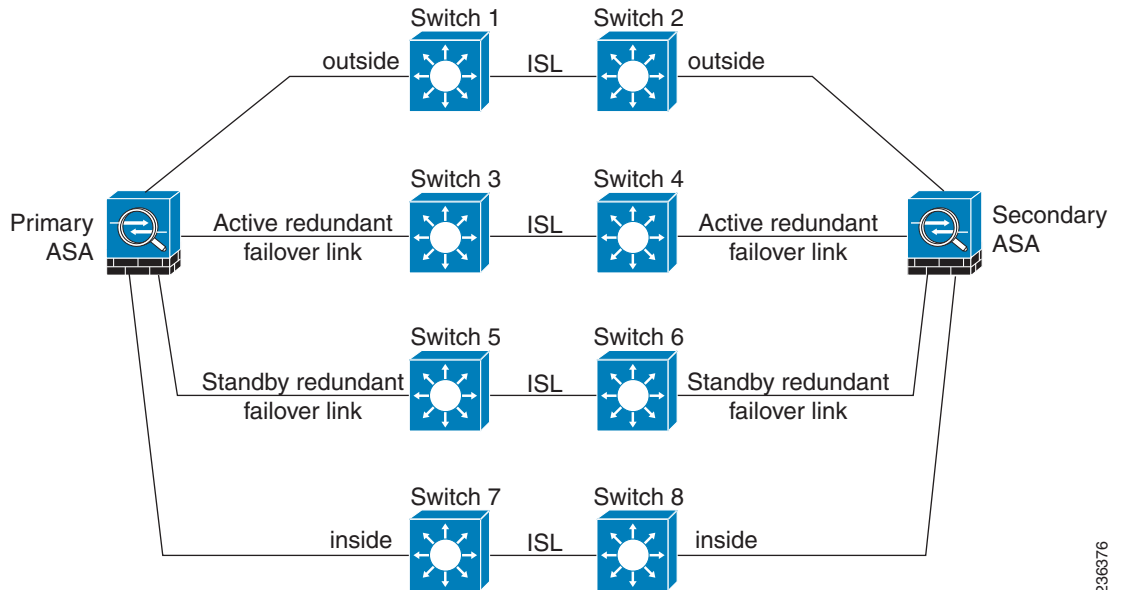
The most reliable failover configurations use a redundant interface on the failover link, as shown in [Figure 7-6](#) and [Figure 7-7](#).

Figure 7-6 Connecting with Redundant Interfaces



236375

Figure 7-7 Connecting with Inter-switch Links



236376

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 6-11](#) for more information. In single context mode, you can manually configure virtual MAC addresses; see the [“Configuring Active/Active Failover” section on page 7-30](#) 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 7-8](#)
- [Inter-Chassis Failover, page 7-9](#)

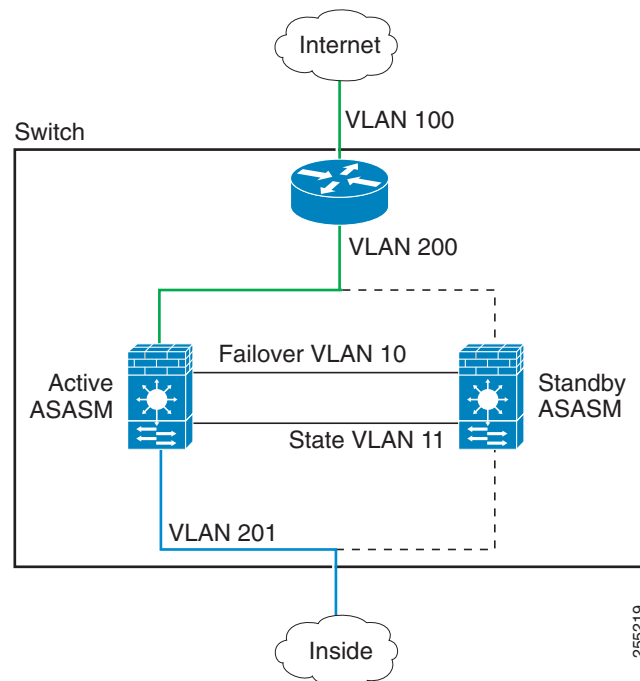
Intra-Chassis Failover

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 7-9](#).

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 7-8 shows a typical intra-switch configuration.

Figure 7-8 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 7-9 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.

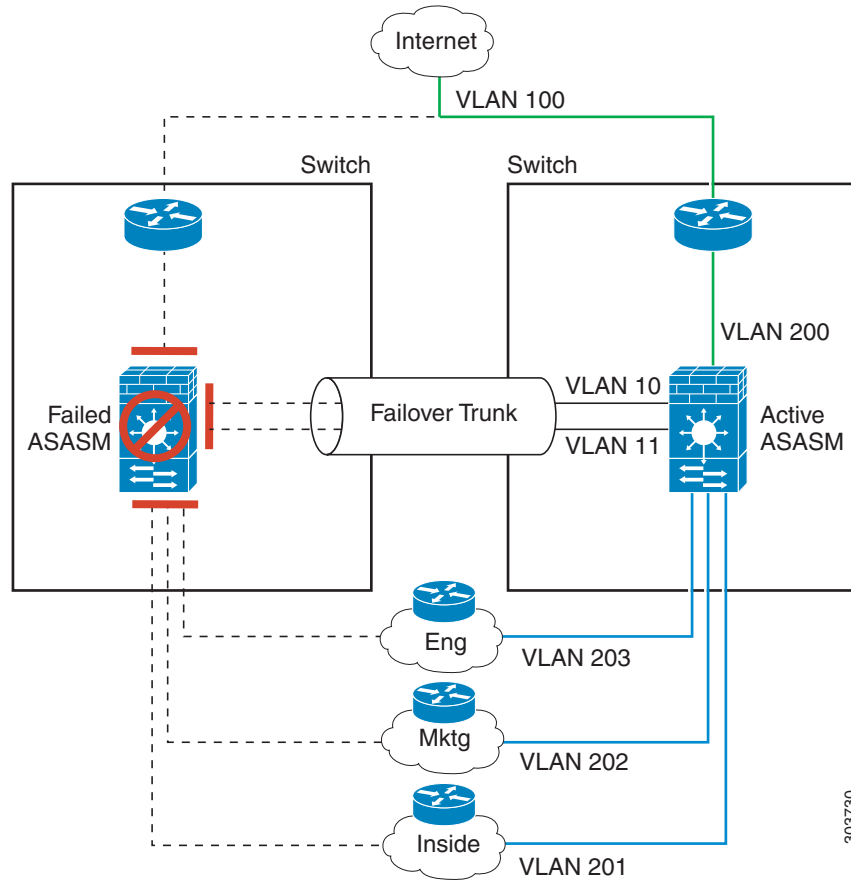
Figure 7-9 Normal Operation



303729

If the primary ASASM fails, then the secondary ASASM becomes active and successfully passes the firewall VLANs (Figure 7-10).

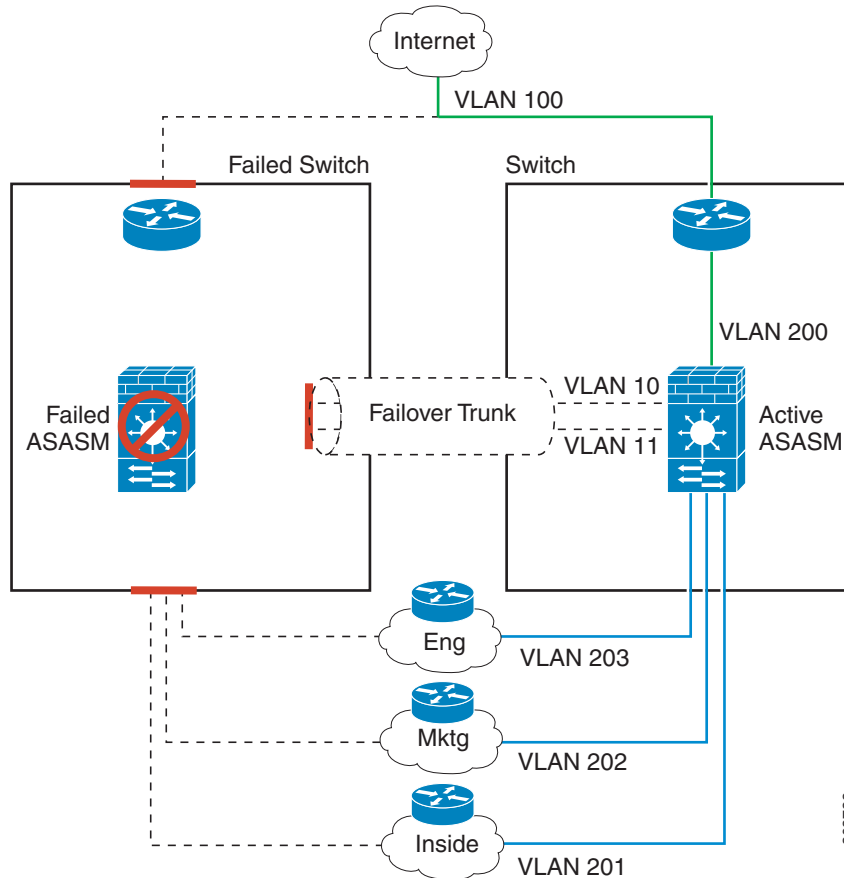
Figure 7-10 ASASM Failure



303730

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 7-11).

Figure 7-11 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 7-13](#)
- [Stateful Failover, page 7-13](#)



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 7-13](#)
- [Unsupported Features, page 7-14](#)

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.

**Note**

Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

- 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 7-15](#)
- [Transparent Mode Requirements for Modules, page 7-15](#)

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.

```
access-list id ethertype deny bpdu
access-group id in interface inside_name
access-group id in interface outside_name
```

Blocking BPDUs disables STP on the switch. Be sure not to have any loops involving the ASA in your network layout.

If neither of the above options are possible, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:

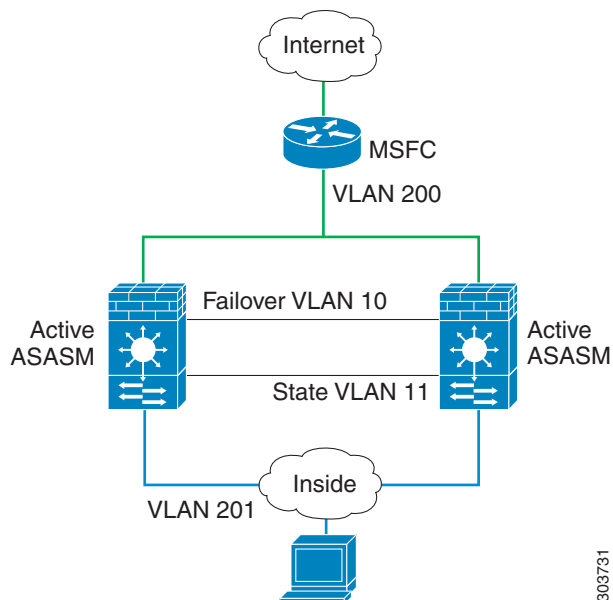
- Disable interface monitoring.
- Increase interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the interface holdtime.

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 7-12](#)). 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 7-12 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 7-16](#)
- [Interface Monitoring, page 7-17](#)

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 7-1 shows the minimum, default, and maximum failover times.

Table 7-1 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 7-18](#)
- [Command Replication, page 7-19](#)

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 according to the [“Saving Configuration Changes” section on page 3-25](#).

**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, commands 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 commands 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 7-20](#)
- [Active Unit Determination at Startup, page 7-20](#)
- [Failover Events, page 7-20](#)

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 7-2 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 7-2 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 7-22](#)
- [Primary/Secondary Roles and Active/Standby Status for a Failover Group, page 7-22](#)
- [Failover Events, page 7-23](#)

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 1 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 7-3 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 7-3 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 7-2.

Guidelines and Limitations

For Auto Update guidelines with failover, see the [“Auto Update Server Support in Failover Configurations”](#) section on page 42-36.

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 7-24 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

- [Configuring the Primary Unit for Active/Standby Failover, page 7-26](#)
- [Configuring the Secondary Unit for Active/Standby Failover, page 7-30](#)

Configuring the Primary Unit for Active/Standby Failover

Follow the steps in this section to configure the primary in an Active/Standby failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit.

Prerequisites

- Configure standby IP addresses for all interfaces except for the failover and state links according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)
- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	<code>failover lan unit primary</code>	Designates this unit as the primary unit.
Step 2	<code>failover lan interface if_name interface_id</code>	Specifies the interface to be used as the failover link. This interface cannot be used for any other purpose (except, optionally, the state link).
	<p>Example:</p> <pre>ciscoasa(config)# failover lan interface folink gigabitethernet0/3</pre>	<p>The <i>if_name</i> argument assigns a name to the interface.</p> <p>The <i>interface_id</i> argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the <i>interface_id</i> specifies a VLAN ID.</p> <p>Note Although you can use an EtherChannel as a failover or state link, 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.</p>

	Command	Purpose
Step 3	<p>failover interface ip <i>failover_if_name</i> {<i>ip_address mask</i> <i>ipv6_address/prefix</i>} standby <i>ip_address</i></p> <p>Example: ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2</p> <p>Or: ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71</p>	<p>Assigns the active and standby IP addresses to the failover link. This address should be on an unused subnet.</p> <p>The standby IP address must be in the same subnet as the active IP address.</p>
Step 4	<p>interface <i>failover_interface_id</i> no shutdown</p> <p>Example: ciscoasa(config)# interface gigabitethernet 0/3 ciscoasa(config-if)# no shutdown</p>	<p>Enables the failover link.</p>
Step 5	<p>failover link <i>if_name interface_id</i></p> <p>Example: ciscoasa(config)# failover link statelink gigabitethernet0/4</p>	<p>(Optional) Specifies the interface you want to use as the state link. We recommend specifying a separate interface from the failover link or data interfaces.</p> <p>The <i>if_name</i> argument assigns a name to the interface.</p> <p>The <i>interface_id</i> argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the <i>interface_id</i> specifies a VLAN ID.</p> <p>Note Although you can use an EtherChannel as a failover or state link, 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.</p>
Step 6	<p>failover interface ip <i>state_if_name</i> {<i>ip_address mask</i> <i>ipv6_address/prefix</i>} standby <i>ip_address</i></p> <p>Example: ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2</p> <p>Or: ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby 2001:a0a:b00:a::a0a:b71</p>	<p>If you specified a separate state link, assigns the active and standby IP addresses to the state link. This address should be on an unused subnet, different from the failover link.</p> <p>The standby IP address must be in the same subnet as the active IP address.</p> <p>Skip this step if you are sharing the state link.</p>

Command	Purpose
<p>Step 7</p> <pre>interface state_interface_id no shutdown</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabitethernet 0/4 ciscoasa(config-if)# no shutdown</pre>	<p>If you specified a separate state link, enables the state link.</p> <p>Skip this step if you are sharing the state link.</p>
<p>Step 8</p> <p>(Optional) Do one of the following to encrypt communications on the failover and state links:</p> <pre>failover ipsec pre-shared-key [0 8] key</pre> <p>Example:</p> <pre>ciscoasa(config)# failover ipsec pre-shared-key a3rynsun</pre>	<p>(Preferred) Establishes IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications. The <i>key</i> can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels.</p> <p>If you use a master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from more system:running-config output), specify that the key is encrypted by using the 8 keyword. 0 is used by default, specifying an unencrypted password.</p> <p>Note The failover ipsec pre-shared-key shows as ***** in show running-config output; this obscured key is not copyable.</p> <p>If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.</p> <p>You cannot use both IPsec encryption and the legacy failover key encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), you must first remove the failover key using the no failover key command before you configure IPsec encryption.</p> <p>Note Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.</p>

Command	Purpose
<p>failover key [0 8] {hex key shared_secret}</p> <p>Example: ciscoasa(config)# failover key johncr1cht0n</p>	<p>(Optional) Encrypts failover communication on the failover and state links using a <i>shared_secret</i>, from 1 to 63 characters, or a 32-character hex key. For the <i>shared_secret</i>, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units.</p> <p>If you use a master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from more system:running-config output), specify that the shared secret or hex key is encrypted by using the 8 keyword. 0 is used by default, specifying an unencrypted password.</p> <p>Note The failover key shared secret shows as ***** in show running-config output; this obscured key is not copyable.</p> <p>If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.</p>
<p>Step 9 failover</p> <p>Example: ciscoasa(config)# failover</p>	<p>Enables failover.</p>
<p>Step 10 write memory</p> <p>Example: ciscoasa(config)# write memory</p>	<p>Saves the system configuration to flash memory.</p>

Examples

The following example configures the failover parameters for the primary unit:

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3
failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
interface gigabitethernet 0/3
    no shutdown
failover link statelink gigabitethernet0/4
failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2
interface gigabitethernet 0/4
    no shutdown
failover ipsec pre-shared-key a3rynsun
failover
```

Configuring the Secondary Unit for Active/Standby Failover

The only configuration required on the secondary unit is for the failover link. The secondary unit requires these commands to communicate initially with the primary unit. After the primary unit sends its configuration to the secondary unit, the only permanent difference between the two configurations is the **failover lan unit** command, which identifies each unit as primary or secondary.

Prerequisites

- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

-
- Step 1** Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. See the “[Configuring the Primary Unit for Active/Standby Failover](#)” section on page 7-26.

For example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
INFO: Non-failover interface config is cleared on GigabitEthernet0/3 and its
sub-interfaces
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby
172.27.48.2
ciscoasa(config)# interface gigabitethernet 0/3
    no shutdown
ciscoasa(config)# failover link statelink gigabitethernet0/4
INFO: Non-failover interface config is cleared on GigabitEthernet0/4 and its
sub-interfaces
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
ciscoasa(config)# interface gigabitethernet 0/4
    no shutdown
ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
ciscoasa(config)# failover
```

- Step 2** After the failover configuration syncs, save the configuration to flash memory:

```
ciscoasa(config)# write memory
```

Configuring Active/Active Failover

- [Configuring the Primary Unit for Active/Active Failover, page 7-31](#)
- [Configuring the Secondary Unit for Active/Active Failover, page 7-35](#)

Configuring the Primary Unit for Active/Active Failover

Follow the steps in this section to configure the primary unit in an Active/Active failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit.

Prerequisites

- Enable multiple context mode according to the [“Enabling or Disabling Multiple Context Mode” section on page 6-16](#).
- Configure standby IP addresses for all interfaces except for the failover and state links according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\)”](#), or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\)”](#).
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	<code>failover lan unit primary</code>	Designates this unit as the primary unit.
Step 2	<p><code>failover lan interface if_name interface_id</code></p> <p>Example: <pre>ciscoasa(config)# failover lan interface folink gigabitethernet0/3</pre></p>	<p>Specifies the interface to be used as the failover link. This interface cannot be used for any other purpose (except, optionally, the state link).</p> <p>The <i>if_name</i> argument assigns a name to the interface.</p> <p>The <i>interface_id</i> argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the <i>interface_id</i> specifies a VLAN ID.</p> <p>Note Although you can use an EtherChannel as a failover or state link, 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.</p>
Step 3	<p><code>failover interface ip if_name</code> <code>{ip_address mask ipv6_address/prefix}</code> <code>standby ip_address</code></p> <p>Example: <pre>ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2</pre></p> <p>Or: <pre>ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71</pre></p>	<p>Assigns the active and standby IP addresses to the failover link. This address should be on an unused subnet.</p> <p>The standby IP address must be in the same subnet as the active IP address.</p>

	Command	Purpose
Step 4	<pre>interface failover_interface_id no shutdown</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabitethernet 0/3 ciscoasa(config-if)# no shutdown</pre>	Enables the failover link.
Step 5	<pre>failover link if_name interface_id</pre> <p>Example:</p> <pre>ciscoasa(config)# failover link statelink gigabitethernet0/4</pre>	<p>(Optional) Specifies the interface you want to use as the state link. We recommend specifying a separate interface from the failover link or data interfaces.</p> <p>The <i>if_name</i> argument assigns a name to the interface.</p> <p>The <i>interface_id</i> argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the <i>interface_id</i> specifies a VLAN ID.</p> <p>Note Although you can use an EtherChannel as a failover or state link, 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.</p>
Step 6	<pre>failover interface ip state_if_name {ip_address mask ipv6_address/prefix} standby ip_address</pre> <p>Example:</p> <pre>ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2</pre> <p>Or:</p> <pre>ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby 2001:a0a:b00:a::a0a:b71</pre>	<p>If you specified a separate state link, assigns the active and standby IP addresses to the state link. This address should be on an unused subnet, different from the failover link.</p> <p>The standby IP address must be in the same subnet as the active IP address.</p> <p>Skip this step if you are sharing the state link.</p>
Step 7	<pre>interface state_interface_id no shutdown</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabitethernet 0/4 ciscoasa(config-if)# no shutdown</pre>	<p>If you specified a separate state link, enables the state link.</p> <p>Skip this step if you are sharing the state link.</p>
Step 8	(Optional) Do one of the following to encrypt communications on the failover and state links:	

Command	Purpose
<p>failover ipsec pre-shared-key [0 8] <i>key</i></p> <p>Example: <pre>ciscoasa(config)# failover ipsec pre-shared-key a3rynsun</pre></p>	<p>(Preferred) Establishes IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications. The <i>key</i> can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels.</p> <p>If you use a master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from more system:running-config output), specify that the key is encrypted by using the 8 keyword. 0 is used by default, specifying an unencrypted password.</p> <p>Note The failover ipsec pre-shared-key shows as ***** in show running-config output; this obscured key is not copyable.</p> <p>If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.</p> <p>You cannot use both IPsec encryption and the legacy failover key encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), you must first remove the failover key using the no failover key command before you configure IPsec encryption.</p> <p>Note Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.</p>
<p>failover key [0 8] {hex key <i>shared_secret</i>}</p> <p>Example: <pre>ciscoasa(config)# failover key johncr1cht0n</pre></p>	<p>(Optional) Encrypts failover communication on the failover and state links using a <i>shared_secret</i>, from 1 to 63 characters, or a 32-character hex key. For the <i>shared_secret</i>, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units.</p> <p>If you use a master passphrase (see the “Configuring the Master Passphrase” section on page 13-8), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from more system:running-config output), specify that the shared secret or hex key is encrypted by using the 8 keyword. 0 is used by default, specifying an unencrypted password.</p> <p>Note The failover key shared secret shows as ***** in show running-config output; this obscured key is not copyable.</p> <p>If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.</p>

	Command	Purpose
Step 9	<pre>failover group 1</pre> <p>Example: <pre>ciscoasa(config)# failover group 1</pre></p>	Creates failover group 1. By default, this group is assigned to the primary unit. Typically, you assign group 1 to the primary unit, and group 2 to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired using the primary or secondary subcommands.
Step 10	<pre>failover group 2 secondary</pre> <p>Example: <pre>ciscoasa(config)# failover group 2 ciscoasa(config-fover-group)# secondary</pre></p>	Creates failover group 2 and assigns it to the secondary unit.
Step 11	<pre>context name join-failover-group {1 2}</pre> <p>Example: <pre>ciscoasa(config)# context Eng ciscoasa(config-ctx)# join-failover-group 2</pre></p>	Enters the context configuration mode for a given context, and assigns the context to a failover group. Repeat this command for each context. Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1; you cannot assign it to group 2.
Step 12	<pre>failover</pre> <p>Example: <pre>ciscoasa(config)# failover</pre></p>	Enables failover.
Step 13	<pre>write memory</pre> <p>Example: <pre>ciscoasa(config)# write memory</pre></p>	Saves the system configuration to flash memory.

Examples

The following example configures the failover parameters for the primary unit:

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3
failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
interface gigabitethernet 0/3
  no shutdown
failover link statelink gigabitethernet0/4
failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2
interface gigabitethernet 0/4
  no shutdown
failover group 1
failover group 2
  secondary
context admin
  join-failover-group 1
failover ipsec pre-shared-key a3rynsun
failover
```

Configuring the Secondary Unit for Active/Active Failover

The only configuration required on the secondary unit is for the failover link. The secondary unit requires these commands to communicate initially with the primary unit. After the primary unit sends its configuration to the secondary unit, the only permanent difference between the two configurations is the **failover lan unit** command, which identifies each unit as primary or secondary.

Prerequisites

- Enable multiple context mode according to the [“Enabling or Disabling Multiple Context Mode” section on page 6-16](#).
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

- Step 1** Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. You also do not need to enter the **failover group** and **join-failover-group** commands, as they are replicated from the primary unit. See the [“Configuring the Primary Unit for Active/Active Failover” section on page 7-31](#).

For example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
INFO: Non-failover interface config is cleared on GigabitEthernet0/3 and its
sub-interfaces
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby
172.27.48.2
ciscoasa(config)# interface gigabitethernet 0/3
no shutdown
ciscoasa(config)# failover link statelink gigabitethernet0/4
INFO: Non-failover interface config is cleared on GigabitEthernet0/4 and its
sub-interfaces
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
ciscoasa(config)# interface gigabitethernet 0/4
no shutdown
ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
ciscoasa(config)# failover
```

- Step 2** After the failover configuration syncs from the primary unit, save the configuration to flash memory:
- ```
ciscoasa(config)# write memory
```
- Step 3** If necessary, force failover group 2 to be active on the secondary unit:
- ```
failover active group 2
```

Configuring Optional Failover Parameters

You can customize failover settings as desired.

- [Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses, page 7-36](#)
- [Configuring Interface Monitoring, page 7-38](#)
- [Configuring Support for Asymmetrically Routed Packets \(Active/Active Mode\), page 7-39](#)

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


See the “Default Settings” section on page 7-25 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.

Prerequisites

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

Detailed Steps

	Command	Purpose
Step 1	<pre>failover polltime [unit] [msec] poll_time [holdtime [msec] time]</pre> <p>Example: ciscoasa(config)# failover polltime unit msec 200 holdtime msec 800</p>	<p>Changes the unit poll and hold times. In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.</p> <p>You cannot enter a holdtime value that is less than 3 times the unit poll time. With a faster poll time, the ASA can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested.</p> <p>If a unit does not hear hello packet on the failover communication interface for one polling period, additional testing occurs through the remaining interfaces. If there is still no response from the peer unit during the hold time, the unit is considered failed and, if the failed unit is the active unit, the standby unit takes over as the active unit.</p>
Step 2	<pre>failover replication rate conns</pre> <p>Example: ciscoasa(config)# failover replication rate 20000</p>	<p>Sets the HTTP replication rate in connections per second, between 8341 and 50000. The default is 50000. In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.</p>
Step 3	<p>(Active/Active mode only)</p> <pre>failover group {1 2}</pre> <p>Example: ciscoasa(config)# failover group 1 ciscoasa(config-fover-group)#</p>	<p>Specifies the failover group you want to customize.</p>

Command	Purpose
<p>Step 4 (Active/Active Mode Only)</p> <pre>preempt [delay]</pre> <p>Example: ciscoasa(config-fover-group)# preempt 1200</p>	<p>Configures failover group preemption for failover group 1. If one unit boots before the other, then both failover groups become active on that unit, despite the primary or secondary setting. This command causes the failover group to become active on the designated unit automatically when that unit becomes available.</p> <p>You can enter an optional <i>delay</i> 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.</p> <p> 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.</p>
<p>Step 5 For Active/Standby mode:</p> <pre>failover replication http</pre> <p>For Active/Active mode:</p> <pre>replication http</pre> <p>Example: ciscoasa(config)# failover replication http</p> <p>Or</p> <pre>ciscoasa(config-fover-group)# replication http</pre>	<p>Enables HTTP state replication. To allow HTTP connections to be included in the state information replication, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information.</p>
<p>Step 6 For Active/Standby mode:</p> <pre>failover interface-policy num[%]</pre> <p>For Active/Active mode:</p> <pre>interface-policy num[%]</pre> <p>Example: ciscoasa (config)# failover interface-policy 20%</p> <p>Or</p> <pre>ciscoasa(config-fover-group)# interface-policy 20%</pre>	<p>Sets the threshold for failover when interfaces fail. By default, one interface failure causes failover.</p> <p>When specifying a specific number of interfaces, the <i>num</i> argument can be from 1 to 250.</p> <p>When specifying a percentage of interfaces, the <i>num</i> argument can be from 1 to 100.</p>

Command	Purpose
<p>Step 7 For Active/Standby mode:</p> <pre>failover polltime interface [msec] time [holdtime time]</pre> <p>For Active/Active mode:</p> <pre>polltime interface [msec] time [holdtime time]</pre> <p>Example:</p> <pre>ciscoasa(config)# failover polltime interface msec 500 holdtime 5</pre> <p>Or</p> <pre>ciscoasa(config-fover-group)# polltime interface msec 500 holdtime 5</pre>	<p>Changes the interface poll and hold times.</p> <p>Valid values for poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.</p> <p>If the interface link is down, interface testing is not conducted and the standby unit could become active in just one interface polling period if the number of failed interfaces meets or exceeds the configured failover criteria.</p>
<p>Step 8 For Active/Standby mode:</p> <pre>failover mac address phy_if active_mac standby_mac</pre> <p>For Active/Active mode:</p> <pre>mac address phy_if active_mac standby_mac</pre> <p>Example:</p> <pre>ciscoasa(config)# failover mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8</pre> <p>Or</p> <pre>ciscoasa(config-fover-group)# mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8</pre>	<p>Configures the virtual MAC address for an interface.</p> <p>The <i>phy_if</i> argument is the physical name of the interface, such as gigabitethernet0/1.</p> <p>The <i>active_mac</i> and <i>standby_mac</i> arguments are MAC addresses 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.</p> <p>The <i>active_mac</i> address is associated with the active IP address for the interface, and the <i>standby_mac</i> is associated with the standby IP address for the interface.</p> <p>You can also set the MAC address using other commands or 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.</p> <p>Use the show interface command to display the MAC address used by an interface.</p>
<p>Step 9 (Active/Active mode only)</p> <p>Repeat this procedure for the other failover group, if desired.</p>	

Configuring Interface Monitoring

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.

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

```
[no] monitor-interface if_name
```

Enables or disables health monitoring for an interface.

Example:

```
ciscoasa(config)# monitor-interface inside
ciscoasa(config)# no monitor-interface eng1
```

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 7-13 shows an example of an asymmetrically routed packet.

Figure 7-13 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 ASA with 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

	Command	Purpose
Step 1	<p>On the primary unit:</p> <pre>interface phy_if</pre> <p>Example: primary/admin(config)# interface gigabitethernet 0/0</p>	Specifies the interface on the primary unit for which you want to allow asymmetrically routed packets.
Step 2	<pre>asr-group num</pre> <p>Example: primary/admin(config-ifc)# asr-group 1</p>	Sets the ASR group number for the interface. Valid values for <i>num</i> range from 1 to 32.
Step 3	<p>On the secondary unit:</p> <pre>interface phy_if</pre> <p>Example: secondary/ctx1(config)# interface gigabitethernet 0/1</p>	Specifies the similar interface on the secondary unit for which you want to allow asymmetrically routed packets.
Step 4	<pre>asr-group num</pre> <p>Example: secondary/ctx1(config-ifc)# asr-group 1</p>	Sets the ASR group number for the interface to match the primary unit interface.

Example

The two units have the following configuration (configurations show only the relevant commands). The device labeled SecAppA in the diagram is the primary unit in the failover pair.

Example 7-1 Primary Unit System Configuration

```
interface GigabitEthernet0/1
  description LAN/STATE Failover Interface
interface GigabitEthernet0/2
  no shutdown
interface GigabitEthernet0/3
  no shutdown
interface GigabitEthernet0/4
  no shutdown
interface GigabitEthernet0/5
  no shutdown
failover
failover lan unit primary
failover lan interface folink GigabitEthernet0/1
failover link folink
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
  primary
failover group 2
  secondary
admin-context SecAppA
context admin
```

```
allocate-interface GigabitEthernet0/2
allocate-interface GigabitEthernet0/3
config-url flash:/admin.cfg
join-failover-group 1
context SecAppB
allocate-interface GigabitEthernet0/4
allocate-interface GigabitEthernet0/5
config-url flash:/ctx1.cfg
join-failover-group 2
```

Example 7-2 SecAppA Context Configuration

```
interface GigabitEthernet0/2
 nameif outsideISP-A
 security-level 0
 ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
 asr-group 1
interface GigabitEthernet0/3
 nameif inside
 security-level 100
 ip address 10.1.0.1 255.255.255.0 standby 10.1.0.11
monitor-interface outside
```

Example 7-3 SecAppB Context Configuration

```
interface GigabitEthernet0/4
 nameif outsideISP-B
 security-level 0
 ip address 192.168.2.2 255.255.255.0 standby 192.168.2.1
 asr-group 1
interface GigabitEthernet0/5
 nameif inside
 security-level 100
 ip address 10.2.20.1 255.255.255.0 standby 10.2.20.11
```

Managing Failover

- [Forcing Failover, page 7-42](#)
- [Disabling Failover, page 7-43](#)
- [Restoring a Failed Unit, page 7-44](#)
- [Re-Syncing the Configuration, page 7-44](#)
- [Testing the Failover Functionality, page 7-44](#)

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

Command	Purpose
<p>For Active/Standby mode on the standby unit: failover active</p> <p>For Active/Active mode on the standby unit: failover active [group <i>group_id</i>]</p> <p>Example: standby# failover active</p> <p>Or: standby# failover active group 1</p>	<p>Forces a failover when entered on the <i>standby</i> unit. The standby unit becomes the active unit.</p> <p>If you specify the group <i>group_id</i>, then this command forces a failover when entered on the <i>standby</i> unit for the specified Active/Active failover group. The standby unit becomes the active unit for the failover group.</p>
<p>For Active/Standby mode on the active unit: no failover active</p> <p>For Active/Active mode on the active unit: no failover active [group <i>group_id</i>]</p> <p>Example: active# no failover active</p> <p>Or: active# no failover active group 1</p>	<p>Forces a failover when entered on the <i>active</i> unit. The active unit becomes the standby unit.</p> <p>If you specify the group <i>group_id</i>, then this command forces a failover when entered on the <i>active</i> unit for the specified failover group. The active unit becomes the standby unit for the failover group.</p>

Disabling Failover

To disable failover, perform the following procedure.

Prerequisites

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

Detailed Steps

Command	Purpose
<p>no failover</p> <p>Example: ciscoasa(config)# no failover</p>	<p>Disables failover.</p> <p>Disabling failover on an Active/Standby pair causes the active and standby state of each unit to be maintained until you reload. For example, the standby unit remains in standby mode so that both units do not start passing traffic. To make the standby unit active (even with failover disabled), see the “Forcing Failover” section on page 7-42.</p> <p>Disabling failover on an Active/Active failover pair causes the failover groups to remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer.</p>

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

Command	Purpose
For Active/Standby mode: <code>failover reset</code>	Restores a failed unit to an unfailed state. Restoring a failed unit to an unfailed state does not automatically make it active; restored units remain in the standby state until made active by failover (forced or natural). An exception is a failover group (Active/Active mode only) configured with failover preemption. If previously active, a failover group becomes active if it is configured with preemption and if the unit on which it failed is the preferred unit.
For Active/Active mode: <code>failover reset [group group_id]</code>	
Example: ciscoasa(config)# failover reset	If you specify the group group_id , this command restores a failed Active/Active failover group to an unfailed state.
Or: ciscoasa(config)# failover reset group 1	

Re-Syncing the Configuration

If you enter the **write standby** command on the active unit, the standby unit clears its running configuration (except for the failover commands used to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

For multiple context mode, when you enter the **write standby** command in the system execution space, all contexts are replicated. If you enter the **write standby** command within a context, the command replicates only the context configuration.

Replicated commands are stored in the running configuration.

Testing the Failover Functionality

To test failover functionality, perform the following procedure.

Detailed Steps

-
- Step 1** Test that your active unit is passing traffic as expected by using FTP (for example) to send a file between hosts on different interfaces.
 - Step 2** Force a failover by entering the following command on the active unit:
 Active/Standby mode:
`ciscoasa(config)# no failover active`
 Active/Active mode:


```
ciscoasa(config)# no failover active group group_id
```

Step 3 Use FTP to send another file between the same two hosts.

Step 4 If the test was not successful, enter the **show failover** command to check the failover status.

Step 5 When you are finished, you can restore the unit to active status by enter the following command on the newly active unit:

Active/Standby mode:

```
ciscoasa(config)# no failover active
```

Active/Active mode:

```
ciscoasa(config)# failover active group group_id
```

**Note**

When an ASA interface goes down, for failover it is still considered to be a unit issue. If the ASA detects that an interface is down, failover occurs immediately, without waiting for the interface holdtime. The interface holdtime is only useful when the ASA considers its status to be OK, although it is not receiving hello packets from the peer. To simulate interface holdtime, shut down the VLAN on the switch to prevent peers from receiving hello packets from each other.

Remote Command Execution

Remote command execution lets you send commands entered at the command line to a specific failover peer.

- [Sending a Command, page 7-45](#)
- [Changing Command Modes, page 7-46](#)
- [Security Considerations, page 7-47](#)
- [Limitations of Remote Command Execution, page 7-47](#)

Sending a Command

Because configuration commands are replicated from the active unit or context to the standby unit or context, you can use the **failover exec** command to enter configuration commands on the correct unit, no matter which unit you are logged in to. For example, if you are logged in to the standby unit, you can use the **failover exec active** command to send configuration changes to the active unit. Those changes are then replicated to the standby unit. Do not use the **failover exec** command to send configuration commands to the standby unit or context; those configuration changes are not replicated to the active unit and the two configurations will no longer be synchronized.

Output from configuration, exec, and **show** commands is displayed in the current terminal session, so you can use the **failover exec** command to issue **show** commands on a peer unit and view the results in the current terminal.

You must have sufficient privileges to execute a command on the local unit to execute the command on the peer unit.

Detailed Steps

Step 1 If you are in multiple context mode, use the **changeto context** *name* command to change to the context you want to configure. You cannot change contexts on the failover peer with the **failover exec** command.

Step 2 Use the following command to send commands to the specified failover unit:

```
ciscoasa(config)# failover exec {active | mate | standby}
```

Use the **active** or **standby** keyword to cause the command to be executed on the specified unit, even if that unit is the current unit. Use the **mate** keyword to cause the command to be executed on the failover peer.

Commands that cause a command mode change do not change the prompt for the current session. You must use the **show failover exec** command to display the command mode the command is executed in. See “[Changing Command Modes](#)” for more information.

Changing Command Modes

The **failover exec** command maintains a command mode state that is separate from the command mode of your terminal session. By default, the **failover exec** command mode starts in global configuration mode for the specified device. You can change that command mode by sending the appropriate command (such as the **interface** command) using the **failover exec** command. The session prompt does not change when you change modes using **failover exec**.

For example, if you are logged in to global configuration mode of the active unit of a failover pair, and you use the **failover exec active** command to change to interface configuration mode, the terminal prompt remains in global configuration mode, but commands entered using **failover exec** are entered in interface configuration mode.

The following examples show the difference between the terminal session mode and the **failover exec** command mode. In the example, the administrator changes the **failover exec** mode on the active unit to interface configuration mode for the interface GigabitEthernet0/1. After that, all commands entered using **failover exec active** are sent to interface configuration mode for interface GigabitEthernet0/1. The administrator then uses **failover exec active** to assign an IP address to that interface. Although the prompt indicates global configuration mode, the **failover exec active** mode is in interface configuration mode.

```
ciscoasa(config)# failover exec active interface GigabitEthernet0/1
ciscoasa(config)# failover exec active ip address 192.168.1.1 255.255.255.0 standby
192.168.1.2
ciscoasa(config)# router rip
ciscoasa(config-router)#
```

Changing command modes for your current session to the device does not affect the command mode used by the **failover exec** command. For example, if you are in interface configuration mode on the active unit, and you have not changed the **failover exec** command mode, the following command would be executed in global configuration mode. The result would be that your session to the device remains in interface configuration mode, while commands entered using **failover exec active** are sent to router configuration mode for the specified routing process.

```
ciscoasa(config-if)# failover exec active router ospf 100
ciscoasa(config-if)#
```

Use the **show failover exec** command to display the command mode on the specified device in which commands sent with the **failover exec** command are executed. The **show failover exec** command takes the same keywords as the **failover exec** command: **active**, **mate**, or **standby**. The **failover exec** mode for each device is tracked separately.

For example, the following is sample output from the **show failover exec** command entered on the standby unit:

```
ciscoasa(config)# failover exec active interface GigabitEthernet0/1
ciscoasa(config)# sh failover exec active
Active unit Failover EXEC is at interface sub-command mode

ciscoasa(config)# sh failover exec standby
Standby unit Failover EXEC is at config mode

ciscoasa(config)# sh failover exec mate
Active unit Failover EXEC is at interface sub-command mode
```

Security Considerations

The **failover exec** command uses the failover link to send commands to and receive the output of the command execution from the peer unit. You should enable encryption on the failover link to prevent eavesdropping or man-in-the-middle attacks.

Limitations of Remote Command Execution

When you use remote commands you face the following limitations:

- If you upgrade one unit using the zero-downtime upgrade procedure and not the other, both units must be running software that supports the **failover exec** command for the command to work.
- Command completion and context help is not available for the commands in the *cmd_string* argument.
- In multiple context mode, you can only send commands to the peer context on the peer unit. To send commands to a different context, you must first change to that context on the unit to which you are logged in.
- You cannot use the following commands with the **failover exec** command:
 - **changeto**
 - **debug (undebug)**
- If the standby unit is in the failed state, it can still receive commands from the **failover exec** command if the failure is due to a service card failure; otherwise, the remote command execution will fail.
- You cannot use the **failover exec** command to switch from privileged EXEC mode to global configuration mode on the failover peer. For example, if the current unit is in privileged EXEC mode, and you enter **failover exec mate configure terminal**, the **show failover exec mate** output will show that the failover exec session is in global configuration mode. However, entering configuration commands for the peer unit using **failover exec** will fail until you enter global configuration mode on the current unit.
- You cannot enter recursive failover exec commands, such as **failover exec mate failover exec mate** command.

- Commands that require user input or confirmation must use the **/nonconfirm** option.

Monitoring Failover

- [Failover Messages](#), page 7-48
- [Monitoring Failover](#), page 7-49

Failover Messages

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

- [Failover Syslog Messages](#), page 7-48
- [Failover Debug Messages](#), page 7-48
- [SNMP Failover Traps](#), page 7-48

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 44, “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 45, “Configuring SNMP”](#) for more information.

Monitoring Failover

To monitor failover, enter one of the following commands:

Command	Purpose
<code>show failover</code>	Displays information about the failover state of the unit.
<code>show failover group</code>	Displays information about the failover state of the failover group. The information displayed is similar to that of the show failover command but limited to the specified group.
<code>show monitor-interface</code>	Displays information about the monitored interface.
<code>show running-config failover</code>	Displays the failover commands in the running configuration.

For more information about the output of the monitoring commands, refer to the command reference.

Feature History for Failover

Table 7-4 lists the release history for this feature.

Table 7-4 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 command: failover key hex .
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. We modified the following command: failover key [0 8] .

Table 7-4 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
IPv6 support for failover added.	8.2(2)	We modified the following commands: failover interface ip , show failover , ipv6 address , show monitor-interface .
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	9.1(2)	<p>Instead of using the proprietary encryption for the failover key (the failover key command), you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption.</p> <p>Note Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.</p> <p>We introduced or modified the following commands: failover ipsec pre-shared-key, show vpn-sessiondb.</p>



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 8-20](#).

- [Information About ASA Clustering](#), page 8-1
- [Licensing Requirements for ASA Clustering](#), page 8-27
- [Prerequisites for ASA Clustering](#), page 8-27
- [Guidelines and Limitations](#), page 8-28
- [Default Settings](#), page 8-32
- [Configuring ASA Clustering](#), page 8-32
- [Managing ASA Cluster Members](#), page 8-53
- [Monitoring the ASA Cluster](#), page 8-59
- [Configuration Examples for ASA Clustering](#), page 8-63
- [Feature History for ASA Clustering](#), page 8-78

Information About ASA Clustering

- [How the ASA Cluster Fits into Your Network](#), page 8-2
- [Performance Scaling Factor](#), page 8-2
- [Cluster Members](#), page 8-2
- [ASA Cluster Interfaces](#), page 8-4
- [Cluster Control Link](#), page 8-6
- [High Availability within the ASA Cluster](#), page 8-9
- [Configuration Replication](#), page 8-11
- [ASA Cluster Management](#), page 8-11
- [Load Balancing Methods](#), page 8-13

- [Inter-Site Clustering, page 8-16](#)
- [How the ASA Cluster Manages Connections, page 8-18](#)
- [ASA Features and Clustering, page 8-20](#)

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 8-27](#) 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 8-6](#).
- Management access to each ASA for configuration and monitoring. See the “[ASA Cluster Management](#)” section on [page 8-11](#).

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 8-13](#).
- [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 8-15](#).
- [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 8-16](#).

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 8-3](#)
- [Bootstrap Configuration, page 8-3](#)
- [Master and Slave Unit Roles, page 8-3](#)
- [Master Unit Election, page 8-3](#)

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 42-1](#).
- (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 8-16](#) 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 8-21](#).

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 8-21](#) 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 8-4](#)
- [Interface Type Mode, page 8-6](#)

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 8-13](#).

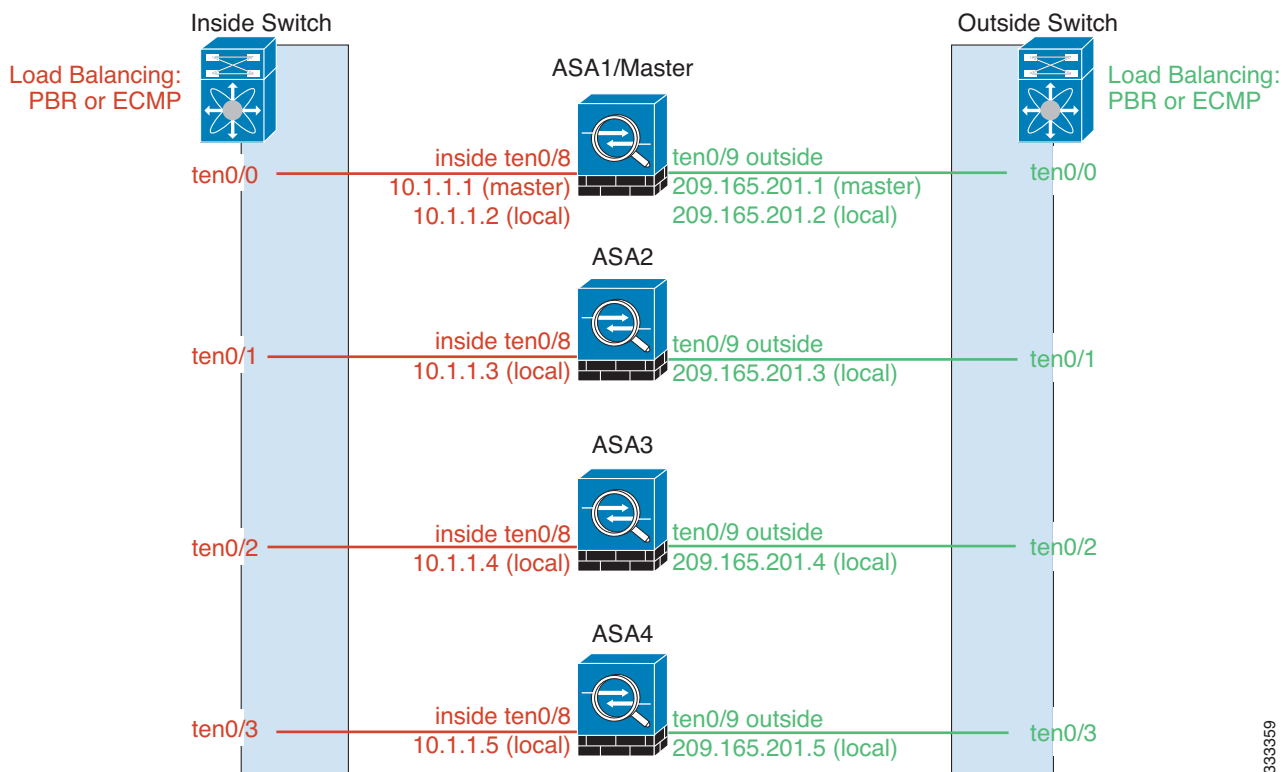


- Individual interfaces (Routed firewall mode only)

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 8-13.



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.



333359

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 8-7](#)
- [Cluster Control Link Interfaces and Network, page 8-7](#)
- [Sizing the Cluster Control Link, page 8-7](#)
- [Cluster Control Link Redundancy, page 8-8](#)

- [Cluster Control Link Reliability, page 8-8](#)
- [Cluster Control Link Failure, page 8-9](#)

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 8-2.)
- Configuration replication. (See the [“Configuration Replication”](#) section on page 8-11.)
- Health monitoring. (See the [“Unit Health Monitoring”](#) section on page 8-9.)

Data traffic includes:

- State replication. (See the [“Data Path Connection State Replication”](#) section on page 8-10.)
- Connection ownership queries and data packet forwarding. (See the [“Rebalancing New TCP Connections Across the Cluster”](#) section on page 8-20.)

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 8-8 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 8-7 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 8-16.

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 8-10](#).



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 8-9](#)
- [Interface monitoring, page 8-9](#)
- [Unit or Interface Failure, page 8-9](#)
- [Data Path Connection State Replication, page 8-10](#)

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 8-10](#).


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 at the console port by entering **cluster name**, and then **enable** (see [Configuring Basic Bootstrap Settings and Enabling Clustering, page 8-45](#)).
- 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 at the console port by entering **cluster name**, and then **enable** (see [Configuring Basic Bootstrap Settings and Enabling Clustering, page 8-45](#)).
- 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 with the **enable** command.

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 8-1](#) for clustering support or lack of support for this kind of traffic.

Table 8-1 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 8-11](#)
- [Management Interface, page 8-11](#)
- [Master Unit Management Vs. Slave Unit Management, page 8-12](#)
- [RSA Key Replication, page 8-12](#)
- [ASDM Connection Certificate IP Address Mismatch, page 8-12](#)

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 8-4.

- [Spanned EtherChannel \(Recommended\), page 8-13](#)
- [Policy-Based Routing \(Routed Firewall Mode Only\), page 8-15](#)
- [Equal-Cost Multi-Path Routing \(Routed Firewall Mode Only\), page 8-16](#)

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 8-13](#)
- [Guidelines for Maximum Throughput, page 8-13](#)
- [Load Balancing, page 8-14](#)
- [EtherChannel Redundancy, page 8-14](#)
- [Connecting to a VSS or vPC, page 8-14](#)

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 9-5.

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 9-32). 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 9-7 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 8-24 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 8-16](#)
- [Sizing the Data Center Interconnect, page 8-17](#)
- [Inter-Site Example, page 8-17](#)

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 8-20](#)); 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 8-18](#)). 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 ($2/2 \times 5$ 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 ($4/2 \times 5$ 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 ($3/2 \times 10$ Gbps).

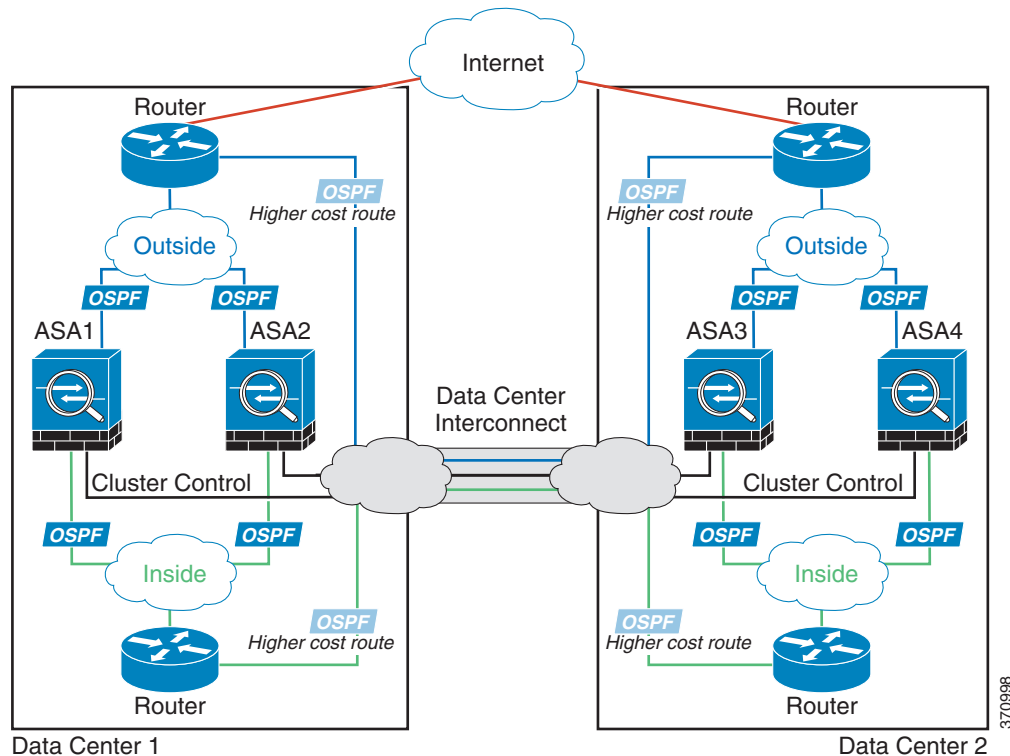
- For 2 members at 2 sites:
 - 2 cluster members total
 - 1 member at each site
 - 10 Gbps cluster control link per member

Reserved DCI bandwidth = 10 Gbps ($1/2 \times 10$ 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 8-18
- [New Connection Ownership](#), page 8-19
- [Sample Data Flow](#), page 8-19
- [Rebalancing New TCP Connections Across the Cluster](#), page 8-20

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.

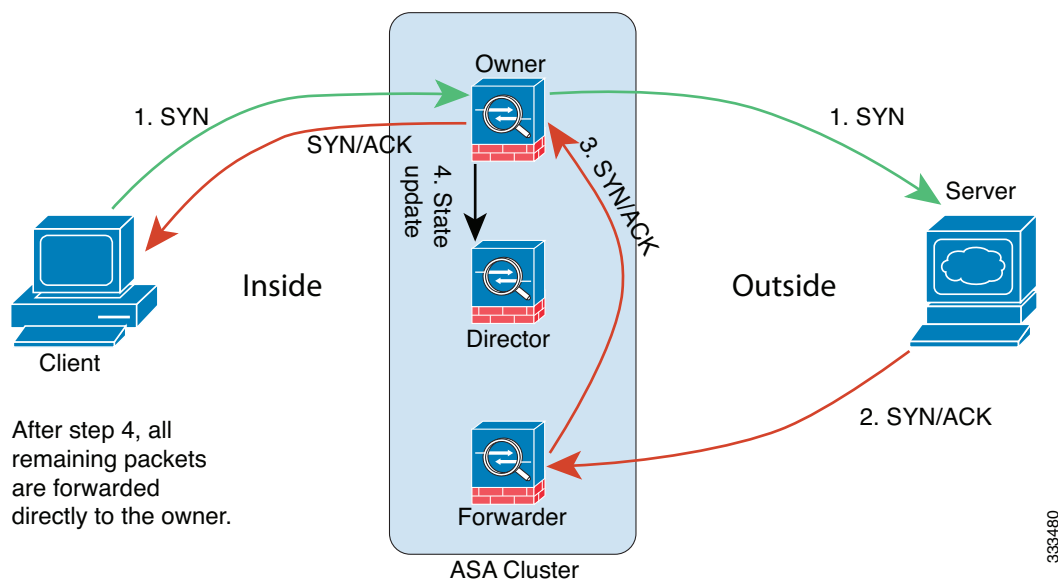
- Forwarder—A unit that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that 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 8-13.

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 8-20](#)
- [Centralized Features, page 8-21](#)
- [Features Applied to Individual Units, page 8-22](#)
- [Dynamic Routing, page 8-22](#)
- [Multicast Routing, page 8-24](#)
- [NAT, page 8-24](#)
- [AAA for Network Access, page 8-25](#)
- [Syslog and Netflow, page 8-26](#)
- [SNMP, page 8-26](#)
- [VPN, page 8-26](#)
- [FTP, page 8-26](#)
- [Cisco TrustSec, page 8-27](#)

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 8-7 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 8-20), 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.

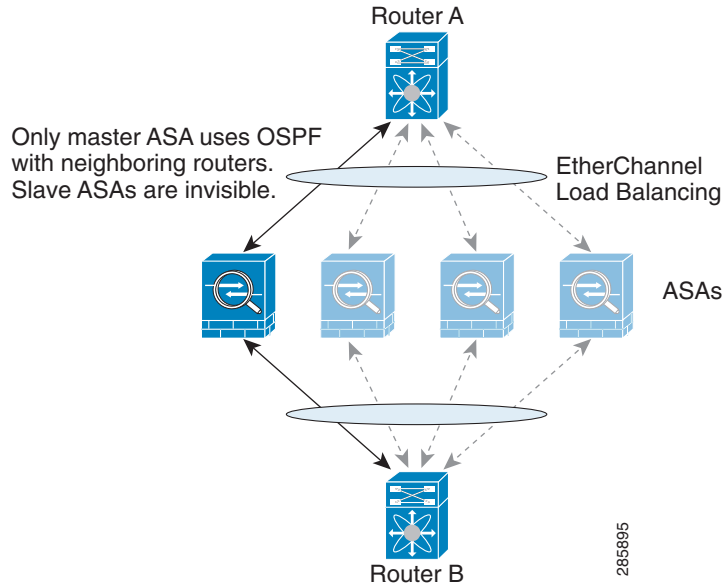
- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster 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 8-22](#)
- [Dynamic Routing in Individual Interface Mode, page 8-23](#)

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.

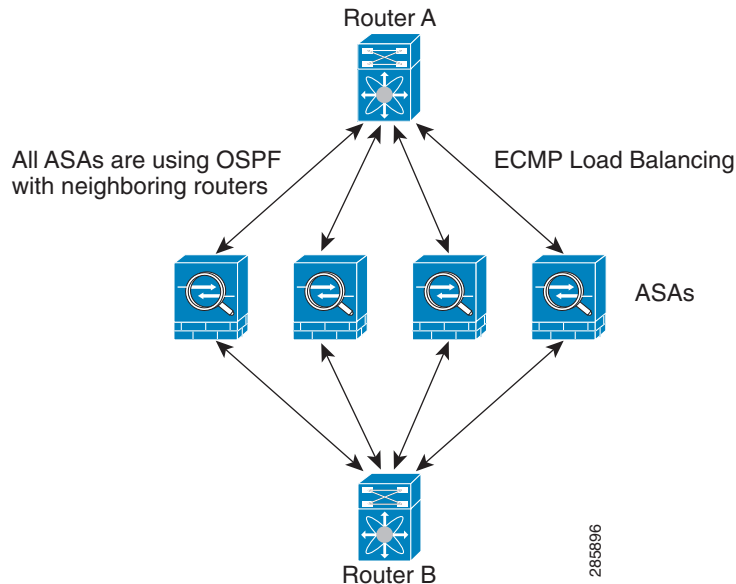
Figure 8-1 Dynamic Routing in Spanned EtherChannel Mode

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 8-2 *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 8-24](#)
- [Multicast Routing in Individual Interface Mode, page 8-24](#)

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 44-18.
- 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 8-2](#) lists supported external hardware and software to interoperate with ASA clustering.

Table 8-2 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 9-35.
- See also the [“ASA Hardware and Software Requirements”](#) section on page 8-3.

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 8-7 for more information.
- ASA 5580
- ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X

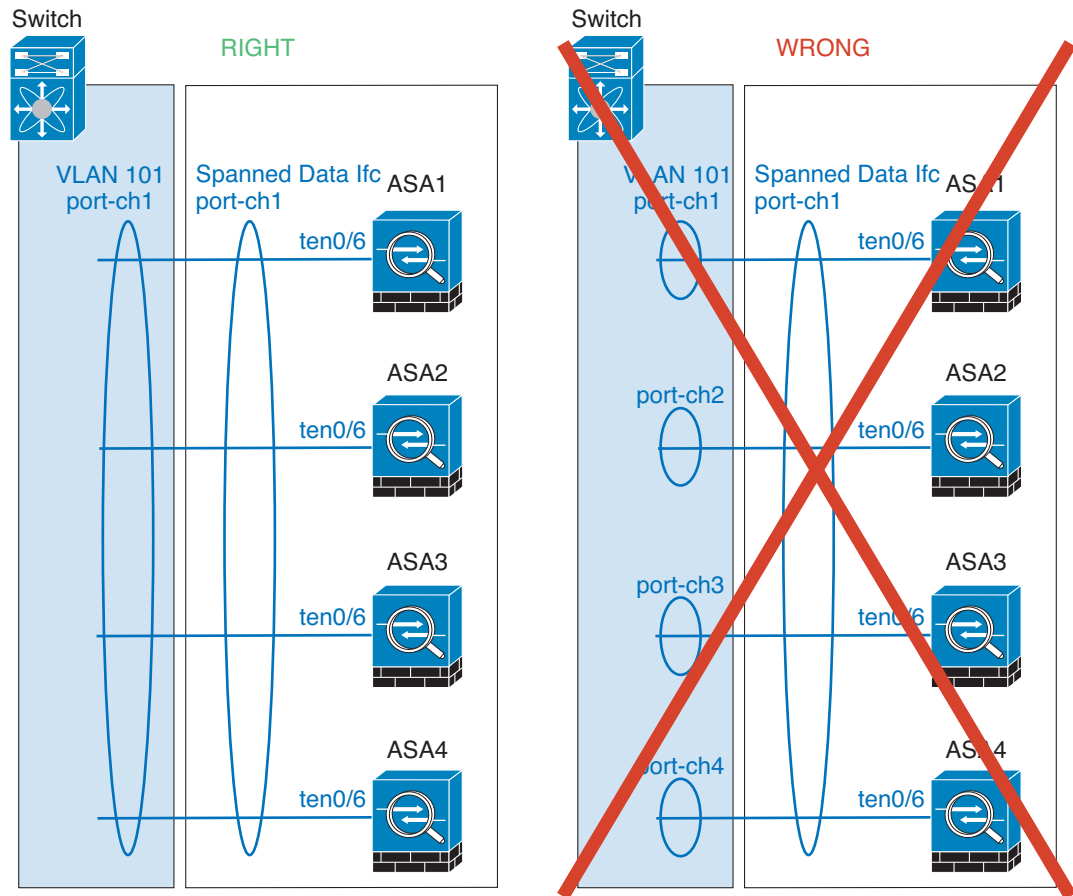
Switch Guidelines

- 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 (in the **port-channel load-balance** command).
- 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.

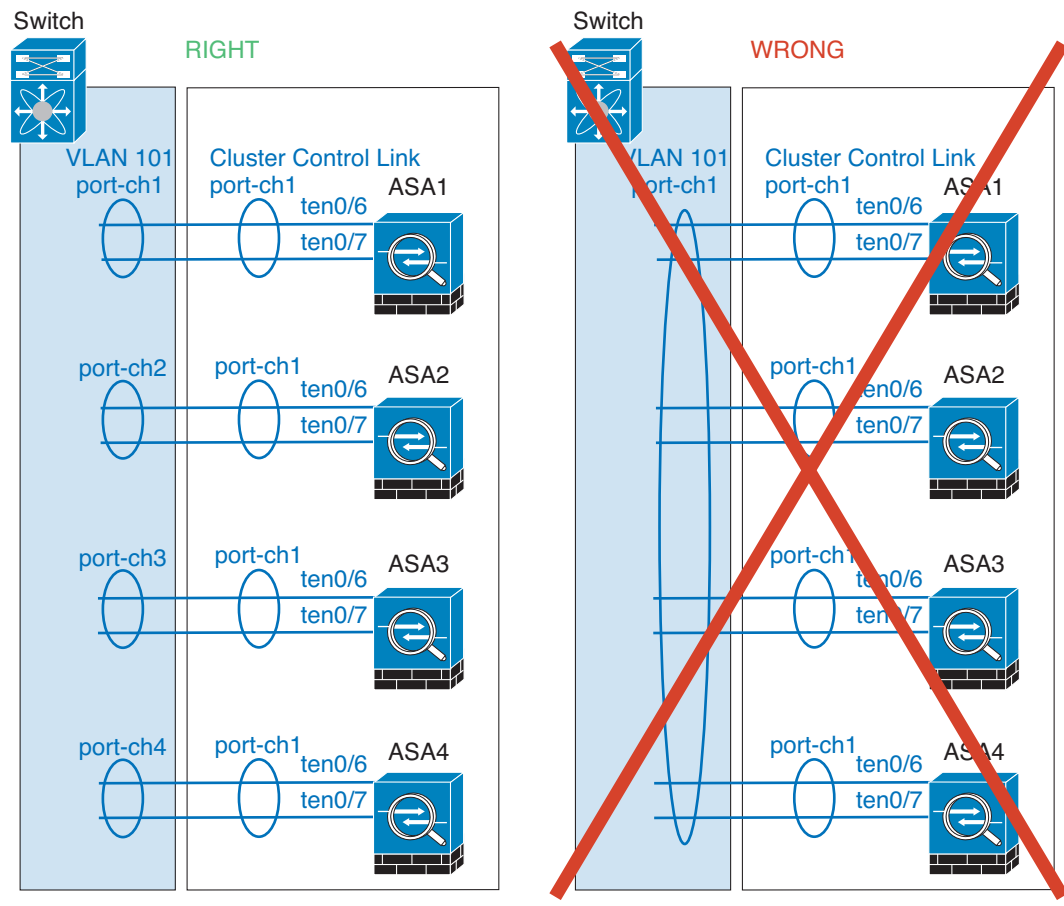
EtherChannel Guidelines

- 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 9-30](#).
- See also the [“EtherChannel Guidelines” section on page 9-13](#).
- 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.



334621

- 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.



333358

Additional Guidelines

- See the [“ASA Hardware and Software Requirements”](#) section on page 8-3.
- For unsupported features with clustering, see the [“Unsupported Features”](#) section on page 8-20.
- 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 8-32
- [Cabling the Cluster Units and Configuring Upstream and Downstream Equipment](#), page 8-33
- [Configuring the Cluster Interface Mode on Each Unit](#), page 8-35
- [Configuring Interfaces on the Master Unit](#), page 8-36
- [Configuring the Master Unit Bootstrap Settings](#), page 8-42
- [Configuring Slave Unit Bootstrap Settings](#), page 8-48

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 8-27.
 - 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 8-33.
 - Step 3** 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 Each Unit”](#) section on page 8-35.
 - Step 4** Configure interfaces for clustering on the master unit. You cannot enable clustering if the interfaces are not cluster-ready. See the [“Configuring Interfaces on the Master Unit”](#) section on page 8-36.
 - Step 5** Configure the bootstrap settings and enable clustering on the master unit. See the [“Configuring the Master Unit Bootstrap Settings”](#) section on page 8-42.
 - Step 6** Configure the bootstrap settings for each slave unit. See the [“Configuring Slave Unit Bootstrap Settings”](#) section on page 8-48.
 - 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 8-20.
-

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

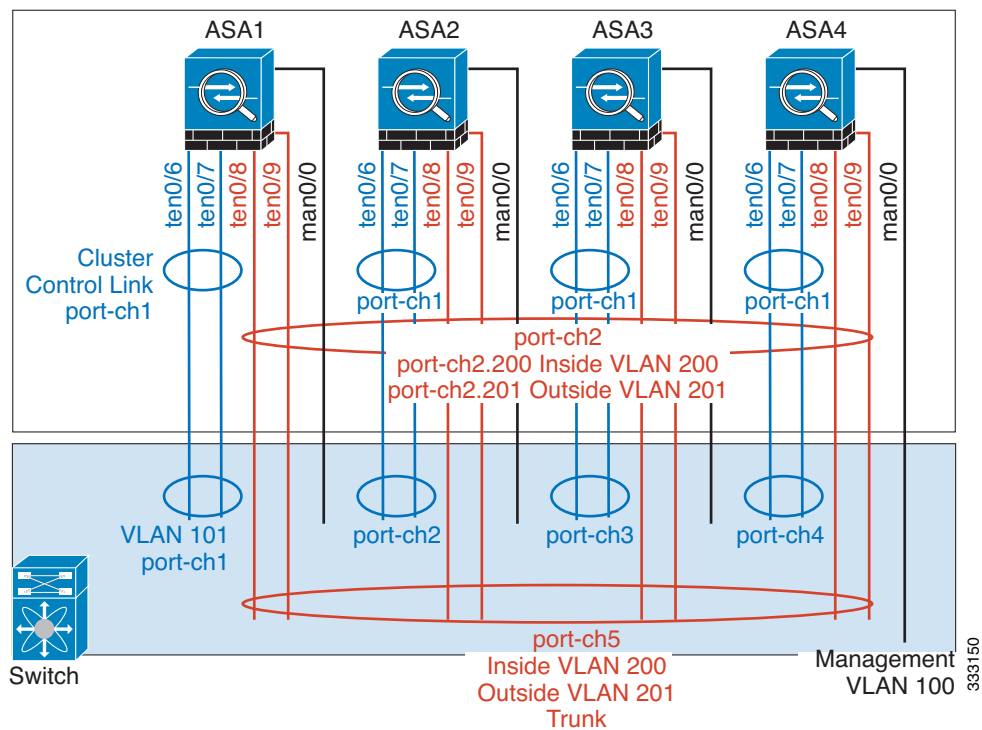
**Note**

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

Configure the cluster interface mode on each unit. See the “[Configuring the Cluster Interface Mode on Each Unit](#)” section on page 8-35.

Configuring the Cluster Interface Mode on Each 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 8-6.

Prerequisites

- You must set the mode separately on each ASA that you want to add to the cluster.
- 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

	Command	Purpose
Step 1	<pre>cluster interface-mode {individual spanned} check-details</pre> <p>Example:</p> <pre>ciscoasa(config)# cluster interface-mode spanned check-details</pre>	<p>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.</p>
Step 2	<pre>cluster interface-mode {individual spanned} force</pre> <p>Example:</p> <pre>ciscoasa(config)# cluster interface-mode spanned force</pre>	<p>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.</p> <p>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.</p> <p>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.</p> <p>To remove the interface mode, enter the no cluster interface-mode command.</p>

What to Do Next

Configure interfaces. See the [“Configuring Interfaces on the Master Unit”](#) section on page 8-36.

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. 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.

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 8-4.

- [Configuring Individual Interfaces \(Recommended for the Management Interface\)](#), page 8-36
- [Configuring Spanned EtherChannels](#), page 8-38

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 8-11 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 Each Unit”](#) section on page 8-35.
- For multiple context mode, perform this procedure in each context. If you are not already in the context configuration mode, enter the **changeto context name** command.
- 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 8-13.
- (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 9-30. 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 9-28. Management-only interfaces cannot be redundant interfaces.
 - For subinterfaces, see the [“Configuring VLAN Subinterfaces and 802.1Q Trunking”](#) section on page 9-33.

Detailed Steps

	Command	Purpose
Step 1	<p>(IPv4)</p> <pre>ip local pool poolname first-address-last-address [mask mask]</pre> <p>(IPv6)</p> <pre>ipv6 local pool poolname ipv6-address/prefix-length number_of_addresses</pre> <p>Example:</p> <pre>ciscoasa(config)# ip local pool ins 192.168.1.2-192.168.1.9 ciscoasa(config-if)# ipv6 local pool insipv6 2001:DB8::1002/32 8</pre>	<p>Configures a pool of Local IP addresses (IPv4 and/or IPv6), one of which will be assigned to each cluster unit for the interface. Include at least as many addresses as there are units in the cluster. If you plan to expand the cluster, include additional addresses. The Main cluster IP address that belongs to the current master unit is <i>not</i> a part of this pool; be sure to reserve an IP address on the same network for the Main cluster IP address.</p> <p>You cannot determine the exact Local address assigned to each unit in advance; to see the address used on each unit, enter the show ip[v6] local pool poolname command. Each cluster member is assigned a member ID when it joins the cluster. The ID determines the Local IP used from the pool.</p>
Step 2	<p>interface interface_id</p> <p>Example:</p> <pre>ciscoasa(config)# interface tengigabitethernet 0/8</pre>	<p>Enters interface configuration mode.</p>
Step 3	<p>(Management interface only)</p> <p>management-only</p> <p>Example:</p> <pre>ciscoasa(config-if)# management-only</pre>	<p>Sets an interface to management-only mode so that it does not pass through traffic.</p> <p>By default, Management type interfaces are configured as management-only. In transparent mode, this command is always enabled for a Management type interface.</p> <p>This setting is required if the cluster interface mode is Spanned.</p>
Step 4	<p>nameif name</p> <p>Example:</p> <pre>ciscoasa(config-if)# nameif inside</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.</p>
Step 5	<p>(IPv4)</p> <pre>ip address ip_address [mask] cluster-pool poolname</pre> <p>(IPv6)</p> <pre>ipv6 address ipv6-address/prefix-length cluster-pool poolname</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0 cluster-pool ins ciscoasa(config-if)# ipv6 address 2001:DB8::1002/32 cluster-pool insipv6</pre>	<p>Sets the Main cluster IP address and identifies the cluster pool. This IP address must be on the same network as the cluster pool addresses, but not be part of the pool. You can configure an IPv4 and/or an IPv6 address.</p> <p>DHCP, PPPoE, and IPv6 autoconfiguration are not supported; you must manually configure the IP addresses.</p>

	Command	Purpose
Step 6	<code>security-level number</code> Example: <code>ciscoasa(config-if)# security-level 100</code>	Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest). See the “ Security Levels ” section on page 11-1.
Step 7	<code>no shutdown</code> Example: <code>ciscoasa(config-if)# no shutdown</code>	Enables the interface.

Examples

The following example configures the Management 0/0 and Management 0/1 interfaces as a device-local EtherChannel, and then configures the EtherChannel as an Individual interface:

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8:45:1002/64 8

interface management 0/0
    channel-group 1 mode active
    no shutdown

interface management 0/1
    channel-group 1 mode active
    no shutdown

interface port-channel 1
    nameif management
    ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
    ipv6 address 2001:DB8:45:1001/64 cluster-pool mgmtipv6
    security-level 100
    management-only
```

What to Do Next

- For spanned interface mode, configure your data interfaces. See the “[Configuring Spanned EtherChannels](#)” section on page 8-38.
- For Individual interface mode, join the cluster. See the “[Configuring the Master Unit Bootstrap Settings](#)” section on page 8-42.

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 Each Unit](#)” section on page 8-35.
- For multiple context mode, start this procedure in the system execution space. If you are not already in the System configuration mode, enter the **changeto system** command.

- For transparent mode, configure the bridge group according to the [“Configuring Bridge Groups” section on page 12-8](#).

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 (The **lACP max-bundle** and **port-channel min-bundle** commands) 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 (see the **port-channel load-balance** command). 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.
- The **lACP port-priority** and **lACP system-priority** commands are not used for a Spanned EtherChannel.
- When using Spanned EtherChannels, the port-channel interface will not come up until clustering is fully enabled (see the [“Configuring the Master Unit Bootstrap Settings” section on page 8-42](#)). 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 9-30](#).
- See also the [“EtherChannel Guidelines” section on page 9-13](#).

Detailed Steps

	Command	Purpose
Step 1	<p>interface <i>physical_interface</i></p> <p>Example: <pre>ciscoasa(config)# interface gigabitethernet 0/0</pre></p>	<p>Specifies the interface you want to add to the channel group, where the <i>physical_interface</i> ID includes the type, slot, and port number as <i>type slot/port</i>. This first interface in the channel group determines the type and speed for all other interfaces in the group.</p>
Step 2	<p>channel-group <i>channel_id</i> mode active [vss-id {1 2}]</p> <p>Example: <pre>ciscoasa(config-if)# channel-group 1 mode active</pre></p>	<p>Assigns this interface to an EtherChannel with the <i>channel_id</i> between 1 and 48. If the port-channel interface for this channel ID does not yet exist in the configuration, one will be added automatically:</p> <p>interface port-channel <i>channel_id</i></p> <p>Only active mode is supported for Spanned EtherChannels.</p> <p>If you are connecting the ASA to two switches in a VSS or vPC, then configure the vss-id keyword to identify to which switch this interface is connected (1 or 2). You must also use the port-channel span-cluster vss-load-balance command for the port-channel interface in Step 6. See also the “Connecting to a VSS or vPC” section on page 8-14 for more information.</p>
Step 3	<p>no shutdown</p> <p>Example: <pre>ciscoasa(config-if)# no shutdown</pre></p>	<p>Enables the interface.</p>
Step 4	<p>(Optional) Add additional interfaces to the EtherChannel by repeating Step 1 through Step 3.</p> <p>Example: <pre>ciscoasa(config)# interface gigabitethernet 0/1 ciscoasa(config-if)# channel-group 1 mode active ciscoasa(config-if)# no shutdown</pre></p>	<p>Multiple interfaces in the EtherChannel per unit are useful for connecting to 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. 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.</p>
Step 5	<p>interface port-channel <i>channel_id</i></p> <p>Example: <pre>ciscoasa(config)# interface port-channel 1</pre></p>	<p>Specifies the port-channel interface. This interface was created automatically when you added an interface to the channel group.</p>
Step 6	<p>port-channel span-cluster [vss-load-balance]</p> <p>Example: <pre>ciscoasa(config-if)# port-channel span-cluster</pre></p>	<p>Sets this EtherChannel as a Spanned EtherChannel.</p> <p>If you are connecting the ASA to two switches in a VSS or vPC, then you should enable VSS load balancing by using the vss-load-balance keyword. This feature ensures that the physical link connections between the ASAs to the VSS (or vPC) pair are balanced. You must configure the vss-id keyword in the channel-group command for each member interface before enabling load balancing (see Step 2).</p>

	Command	Purpose
Step 7	<p>(Optional)</p> <p>You can set the Ethernet properties for the port-channel interface to override the properties set on the Individual interfaces.</p>	<p>This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group. See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section on page 9-26 for Ethernet commands.</p>
Step 8	<p>(Optional)</p> <p>If you are creating VLAN subinterfaces on this EtherChannel, do so now. The rest of this procedure applies to the subinterfaces.</p> <p>Example:</p> <pre>ciscoasa(config)# interface port-channel 1.10 ciscoasa(config-if)# vlan 10</pre>	<p>See the “Configuring VLAN Subinterfaces and 802.1Q Trunking” section on page 9-33.</p>
Step 9	<p>(Multiple Context Mode)</p> <p>Allocate the interface to a context. See the “Configuring a Security Context” section on page 6-19.</p> <p>Then enter:</p> <pre>changeto context name interface port-channel channel_id</pre> <p>Example:</p> <pre>ciscoasa(config)# context admin ciscoasa(config)# allocate-interface port-channell ciscoasa(config)# changeto context admin ciscoasa(config-if)# interface port-channel 1</pre>	<p>For multiple context mode, the rest of the interface configuration occurs within each context.</p>
Step 10	<p>nameif <i>name</i></p> <p>Example:</p> <pre>ciscoasa(config-if)# nameif inside</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.</p>
Step 11	<p>Perform one of the following, depending on the firewall mode:</p> <p>Routed Mode:</p> <p>(IPv4)</p> <pre>ip address ip_address [mask]</pre> <p>(IPv6)</p> <pre>ipv6 address ipv6-prefix/prefix-length</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0 ciscoasa(config-if)# ipv6 address 2001:DB8::1001/32</pre>	<p>Sets the IPv4 and/or IPv6 address. DHCP, PPPoE, and IPv6 autoconfig are not supported.</p>

Command	Purpose
Transparent Mode: bridge-group <i>number</i> Example: ciscoasa(config-if)# bridge-group 1	Assigns the interface to a bridge group, where <i>number</i> is an integer between 1 and 100. You can assign up to four interfaces to a bridge group. You cannot assign the same interface to more than one bridge group. Note that the BVI configuration includes the IP address.
Step 12 security-level <i>number</i> Example: ciscoasa(config-if)# security-level 50	Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest). See the “ Security Levels ” section on page 11-1 .
Step 13 mac-address <i>mac_address</i> Example: ciscoasa(config-if)# mac-address 000C.F142.4CDE	<p>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.</p> <p>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.</p> <p>The <i>mac_address</i> is in H.H.H format, 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.</p> <p>The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.</p>

What to Do Next

Configure the master unit bootstrap settings. See the “[Configuring the Master Unit Bootstrap Settings](#)” section on [page 8-42](#).

Configuring the Master Unit Bootstrap Settings

Each unit in the cluster requires a bootstrap configuration to join the cluster. Typically, the first unit you configure to join the cluster will be the master unit. After you enable clustering, after an election period, the cluster elects a master unit. With only one unit in the cluster initially, that unit will become the master unit. Subsequent units that you add to the cluster will be slave units.

- [Prerequisites, page 8-42](#)
- [Enabling the Cluster Control Link Interface, page 8-43](#)
- [Configuring Basic Bootstrap Settings and Enabling Clustering, page 8-45](#)
- [Configuring Advanced Clustering Settings, page 8-47](#)
- [Examples, page 8-48](#)

Prerequisites

- You must use the console port to enable or disable clustering. You cannot use Telnet or SSH.
- Back up your configurations in case you later want to leave the cluster, and need to restore your configuration.
- For multiple context mode, complete these procedures in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.
- We recommend enabling jumbo frame reservation for use with the cluster control link. See the [“Enabling Jumbo Frame Support \(Supported Models\)”](#) section on page 9-35.
- With the exception of the cluster control link, any interfaces in your configuration must be configured with a cluster IP pool or as a Spanned EtherChannel before you enable clustering, depending on your interface mode. If you have pre-existing interface configuration, you can either clear the interface configuration (**clear configure interface**), or convert your interfaces to cluster interfaces according to the [“Configuring Interfaces on the Master Unit”](#) section on page 8-36 before you enable clustering.
- When you add a unit to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Enabling the Cluster Control Link Interface

You need to enable the cluster control link interface before you join the cluster. You will later identify this interface as the cluster control link when you enable clustering.

We recommend that you combine multiple cluster control link interfaces into an EtherChannel if you have enough interfaces. The EtherChannel is local to the ASA, and is not a Spanned EtherChannel.

The cluster control link interface configuration is not replicated from the master unit to slave units; however, you must use the same configuration on each unit. Because this configuration is not replicated, you must configure the cluster control link interfaces separately on each unit.

Prerequisites

Determine the size of the cluster control link by referring to the [“Sizing the Cluster Control Link”](#) section on page 8-7.

Restrictions

- You cannot use a VLAN subinterface as the cluster control link.
- You cannot use a Management *x/x* interface as the cluster control link, either alone or as an EtherChannel.
- For the ASA 5585-X with an ASA IPS module, you cannot use the module interfaces for the cluster control link.

Detailed Steps—Single Interface

	Command	Purpose
Step 1	interface <i>interface_id</i> Example: ciscoasa(config)# interface tengigabitethernet 0/6	Enters interface configuration mode.
Step 2	no shutdown Example: ciscoasa(config-if)# no shutdown	Enables the interface. You only need to enable the interface; do not configure a name for the interface, or any other parameters.

Detailed Steps—EtherChannel Interface

	Command	Purpose
Step 1	interface <i>interface_id</i> Example: ciscoasa(config)# interface tengigabitethernet 0/6	Enters interface configuration mode.
Step 2	channel-group <i>channel_id</i> mode on Example: ciscoasa(config-if)# channel-group 1 mode on	Assigns this physical interface to an EtherChannel with the <i>channel_id</i> between 1 and 48. If the port-channel interface for this channel ID does not yet exist in the configuration, one will be added automatically: interface port-channel <i>channel_id</i> We recommend using the On mode for cluster control link member interfaces to reduce unnecessary traffic on the cluster control link. The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network. Note: We recommend setting <i>data</i> EtherChannels to Active mode.
Step 3	no shutdown Example: ciscoasa(config-if)# no shutdown	Enables the interface.
Step 4	interface <i>interface_id</i> channel-group <i>channel_id</i> mode on no shutdown Example: ciscoasa(config)# interface tengigabitethernet 0/7 ciscoasa(config-if)# channel-group 1 mode on ciscoasa(config-if)# no shutdown	Repeat for each additional interface you want to add to the EtherChannel.

What to Do Next

Configure the master unit bootstrap settings. See the [Configuring Basic Bootstrap Settings and Enabling Clustering, page 8-45](#).

Configuring Basic Bootstrap Settings and Enabling Clustering

Perform the following steps to configure basic bootstrap settings and to enable clustering.

Detailed Steps

	Command	Purpose
Step 1	<p>(Optional)</p> <p><code>mtu cluster bytes</code></p> <p>Example: <code>ciscoasa(config)# mtu cluster 9000</code></p>	<p>Specifies the maximum transmission unit for the cluster control link interface, between 64 and 65,535 bytes. The default MTU is 1500 bytes.</p> <p>Note We suggest setting the MTU to 1600 bytes or greater, which requires you to enable jumbo frame reservation before continuing with this procedure. See the “Enabling Jumbo Frame Support (Supported Models)” section on page 9-35. Jumbo frame reservation requires a reload of the ASA.</p> <p>This command is a global configuration command, but is also part of the bootstrap configuration that is not replicated between units.</p>
Step 2	<p><code>cluster group name</code></p> <p>Example: <code>ciscoasa(config)# cluster group pod1</code></p>	<p>Names the cluster and enters cluster configuration mode. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster group per unit. All members of the cluster must use the same name.</p>
Step 3	<p><code>local-unit unit_name</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# local-unit unit1</code></p>	<p>Names this member of the cluster with a unique ASCII string from 1 to 38 characters. Each unit must have a unique name. A unit with a duplicated name will be not be allowed in the cluster.</p>
Step 4	<p><code>cluster-interface interface_id ip ip_address mask</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# cluster-interface port-channel2 ip 192.168.1.1 255.255.255.0</code> INFO: Non-cluster interface config is cleared on Port-Channel2</p>	<p>Specifies the cluster control link interface, preferably an EtherChannel. Subinterfaces and Management interfaces are not allowed. See the “Configuring Interfaces on the Master Unit” section on page 8-36</p> <p>Specify an IPv4 address for the IP address; IPv6 is not supported for this interface. This interface cannot have a <code>nameif</code> configured.</p> <p>For each unit, specify a different IP address on the same network.</p>
Step 5	<p><code>priority priority_number</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# priority 1</code></p>	<p>Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority. See the “Master Unit Election” section on page 8-3 for more information.</p>

Command	Purpose
<p>Step 6 (Optional)</p> <pre>key shared_secret</pre> <p>Example: ciscoasa(cfg-cluster)# key chuntheunavoidable </p>	<p>Sets an authentication 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 key. This command does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.</p>
<p>Step 7 (Optional)</p> <pre>clacp system-mac {mac_address auto} [system-priority number]</pre> <p>Example: ciscoasa(cfg-cluster)# clacp system-mac 000a.0000.aaaa </p>	<p>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. One parameter in cLACP negotiation is a system ID, which is in the format of a MAC address. All ASAs in the cluster use the same system ID: auto-generated by the master unit (the default) and replicated to all slaves; or manually specified in this command in the form <i>H.H.H</i>, where H is a 16-bit hexadecimal digit. (For example, the MAC address 00-0A-00-00-AA-AA is entered as 000A.0000.AAAA.) You might want to manually configure the MAC address for troubleshooting purposes, for example, so you can use an easily identified MAC address. Typically, you would use the auto-generated MAC address.</p> <p>The system priority, between 1 and 65535, 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.</p> <p>This command 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.</p>
<p>Step 8 <code>enable [noconfirm]</code></p> <p>Example: ciscoasa(cfg-cluster)# enable INFO: Clustering is not compatible with following commands: policy-map global_policy class inspection_default inspect skinny policy-map global_policy class inspection_default inspect sip Would you like to remove these commands? [Y]es/[N]o:Y</p> <p>INFO: Removing incompatible commands from running configuration... Cryptochecksum (changed): f16b7fc2 a742727e e40bc0b0 cd169999 INFO: Done</p>	<p>Enables clustering. When you enter the enable command, 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. You are prompted to delete the incompatible commands. If you respond No, then clustering is not enabled. Use the noconfirm keyword to bypass the confirmation and delete incompatible commands automatically.</p> <p>For the first unit enabled, a master unit election occurs. Because the first unit should be the only member of the cluster so far, it will become the master unit. Do not perform any configuration changes during this period.</p> <p>To disable clustering, enter the no enable command.</p> <p>Note If you disable clustering, all data interfaces are shut down, and only the management-only interface is active. If you want to remove the unit from the cluster entirely (and thus want to have active data interfaces), see the “Leaving the Cluster” section on page 8-56.</p>

What to Do Next

- Configure advanced settings. See the “Configuring Advanced Clustering Settings” section on page 8-47.
- Add slave units. See the “Configuring Slave Unit Bootstrap Settings” section on page 8-48.

Configuring Advanced Clustering Settings

Perform the following steps to customize your clustering configuration.

Detailed Steps

	Command	Purpose
Step 1	<p>health-check [holdtime <i>timeout</i>]</p> <p>Example: <pre>ciscoasa(cfg-cluster)# health-check holdtime 5</pre></p>	<p>Customizes the cluster health check feature, which includes unit health monitoring and interface health monitoring. The holdtime 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 <i>unit</i> health check; for interface health, the ASA uses the interface status (up or down).</p> <p>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.</p> <p>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. 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 8-9.</p> <p>Health check is enabled by default. You can disable it using the no form of this command.</p> <p>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, or 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.</p>
Step 2	<p>conn-rebalance [frequency <i>seconds</i>]</p> <p>Example: <pre>ciscoasa(cfg-cluster)# conn-rebalance frequency 60</pre></p>	<p>Enables connection rebalancing for TCP traffic. This command is disabled by default. If enabled, ASAs 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. The default is 5 seconds.</p> <p>Note</p>

	Command	Purpose
Step 3	<code>console-replicate</code>	Enables console replication from slave units to the master unit. This feature is disabled by default. The ASA prints 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.
	Example: <code>ciscoasa(cfg-cluster)# console-replicate</code>	

What to Do Next

Add slave units. See the [“Configuring Slave Unit Bootstrap Settings”](#) section on page 8-48.

Examples

The following example configures a management interface, configures a device-local EtherChannel for the cluster control link, and then enables clustering for the ASA called “unit1,” which will become the master unit because it is added to the cluster first:

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/32 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 0/6
  channel-group 1 mode on
  no shutdown

interface tengigabitethernet 0/7
  channel-group 1 mode on
  no shutdown

cluster group pod1
  local-unit unit1
  cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
  priority 1
  key chuntheunavoidable
  enable noconfirm
```

Configuring Slave Unit Bootstrap Settings

Perform the following procedures to configure the slave units.

- [Prerequisites, page 8-49](#)
- [Enabling the Cluster Control Link Interface, page 8-49](#)
- [Configuring Bootstrap Settings and Joining the Cluster, page 8-50](#)
- [Examples, page 8-52](#)

Prerequisites

- You must use the console port to enable or disable clustering. You cannot use Telnet or SSH.
- Back up your configurations in case you later want to leave the cluster, and need to restore your configuration.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.
- We recommend enabling jumbo frame reservation for use with the cluster control link. See the [“Enabling Jumbo Frame Support \(Supported Models\)”](#) section on page 9-35.
- If you have any interfaces in your configuration that have not been configured for clustering (for example, the default configuration Management 0/0 interface), you can join the cluster as a slave unit (with no possibility of becoming the master in a current election).
- When you add a unit to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Enabling the Cluster Control Link Interface

Configure the same cluster control link interface as you configured for the master unit. See the [“Enabling the Cluster Control Link Interface”](#) section on page 8-43.

Detailed Steps—Single Interface

	Command	Purpose
Step 1	interface <i>interface_id</i> Example: ciscoasa(config)# interface tengigabitethernet 0/6	Enters interface configuration mode.
Step 2	no shutdown Example: ciscoasa(config-if)# no shutdown	Enables the interface. You only need to enable the interface; do not configure a name for the interface, or any other parameters.

Detailed Steps—EtherChannel Interface

	Command	Purpose
Step 1	<code>interface interface_id</code> Example: ciscoasa(config)# interface tengigabitethernet 0/6	Enters interface configuration mode.
Step 2	<code>channel-group channel_id mode on</code> Example: ciscoasa(config-if)# channel-group 1 mode on	Assigns this physical interface to an EtherChannel with the <i>channel_id</i> between 1 and 48. If the port-channel interface for this channel ID does not yet exist in the configuration, one will be added automatically: interface port-channel channel_id We recommend using the On mode for cluster control link member interfaces to reduce unnecessary traffic on the cluster control link. The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network. Note: We recommend setting <i>data</i> EtherChannels to Active mode.
Step 3	<code>no shutdown</code> Example: ciscoasa(config-if)# no shutdown	Enables the interface.
Step 4	<code>interface interface_id</code> <code>channel-group channel_id mode on</code> <code>no shutdown</code> Example: ciscoasa(config)# interface tengigabitethernet 0/7 ciscoasa(config-if)# channel-group 1 mode on ciscoasa(config-if)# no shutdown	Repeat for each additional interface you want to add to the EtherChannel.

What to Do Next

Configure the slave unit bootstrap settings. See the [Configuring Bootstrap Settings and Joining the Cluster, page 8-50](#).

Configuring Bootstrap Settings and Joining the Cluster

Perform the following steps to configure bootstrap settings and join the cluster as a slave unit.

Detailed Steps

	Command	Purpose
Step 1	<p>(Optional)</p> <p><code>mtu cluster bytes</code></p> <p>Example: <code>ciscoasa(config)# mtu cluster 9000</code></p>	<p>Specifies the same MTU that you configured for the master unit.</p> <p>Note We suggest setting the MTU to 1600 bytes or greater, which requires you to enable jumbo frame reservation before continuing with this procedure. See the “Enabling Jumbo Frame Support (Supported Models)” section on page 9-35. Jumbo frame reservation requires a reload of the ASA.</p>
Step 2	<p><code>cluster group name</code></p> <p>Example: <code>ciscoasa(config)# cluster group pod1</code></p>	<p>Identifies the same cluster name that you configured for the master unit.</p>
Step 3	<p><code>local-unit unit_name</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# local-unit unit1</code></p>	<p>Names this member of the cluster with a unique ASCII string from 1 to 38 characters. Each unit must have a unique name. A unit with a duplicated name will not be allowed in the cluster.</p>
Step 4	<p><code>cluster-interface interface_id ip ip_address mask</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# cluster-interface port-channel2 ip 192.168.1.2 255.255.255.0</code> INFO: Non-cluster interface config is cleared on Port-Channel2</p>	<p>Specifies the same cluster control link interface that you configured for the master unit.</p> <p>Specify an IPv4 address for the IP address; IPv6 is not supported for this interface. This interface cannot have a nameif configured.</p> <p>For each unit, specify a different IP address on the same network.</p>
Step 5	<p><code>priority priority_number</code></p> <p>Example: <code>ciscoasa(cfg-cluster)# priority 2</code></p>	<p>Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority. See the “Master Unit Election” section on page 8-3 for more information.</p>

	Command	Purpose
Step 6	(Optional) <code>key shared_secret</code> Example: <code>ciscoasa(cfg-cluster)# key chuntheunavoidable</code>	Sets the same authentication key that you set for the master unit.
Step 7	<code>enable as-slave</code> Example: <code>ciscoasa(cfg-cluster)# enable as-slave</code>	Enables clustering. You can avoid any configuration incompatibilities (primarily the existence of any interfaces not yet configured for clustering) by using the enable as-slave command. This command ensures the slave joins the cluster with no possibility of becoming the master in any current election. Its configuration is overwritten with the one synced from the master unit. To disable clustering, enter the no enable command. Note If you disable clustering, all data interfaces are shut down, and only the management interface is active. If you want to remove the unit from the cluster entirely (and thus want to have active data interfaces), see the “Leaving the Cluster” section on page 8-56.

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 8-20.

Examples

The following example includes the configuration for a slave unit, unit2:

```
interface tengigabitethernet 0/6
  channel-group 1 mode on
  no shutdown

interface tengigabitethernet 0/7
  channel-group 1 mode on
  no shutdown

cluster group pod1
  local-unit unit2
  cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
  priority 2
  key chuntheunavoidable
  enable as-slave
```

370827



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

- (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.

Managing ASA Cluster Members

- [Becoming an Inactive Member, page 8-54](#)
- [Inactivating a Member, page 8-55](#)
- [Leaving the Cluster, page 8-56](#)
- [Changing the Master Unit, page 8-57](#)
- [Executing a Command Cluster-Wide, page 8-58](#)

Configuration > Device Management > High Availability and Scalability > ASA

Configure cluster settings to enable this ASA to participate in a cluster. The name of the cluster and the management interface configuration must be the same for all devices in the cluster.

i Participating in ASA cluster will disable failover and VPN load balancing features.

Participate in ASA cluster

Configure ASA cluster settings

Cluster Name:

Member Name:

Member Priority: (1 - 100)

Shared Key: (optional)

Enable connection rebalancing across all the ASA in the cluster

Interval Between Connection Rebalancing: seconds (1 - 360)

Enable health monitoring of this device within the cluster

Time to Wait Before Device Considered Failed: seconds (0.8 - 45)

Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support

Replicate console output to the master's console

Cluster Control Link

Interface:

IP Address: Subnet Mask:

MTU: (64 - 65535)

Cluster LACP

Virtual System MAC Address: Auto-generate

Specify:

System Priority: (1 - 65535)

370826

- (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.

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 8-56](#). 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

- You must use the console port; you cannot enable or disable clustering from a remote CLI connection.
- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	<code>cluster group name</code> Example: <code>ciscoasa(config)# cluster group pod1</code>	Enters cluster configuration mode.
Step 2	<code>no enable</code> Example: <code>ciscoasa(cfg-cluster)# no enable</code>	Disables clustering. 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 Member

To inactivate a member from any unit, 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 8-56](#). 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, enter the **changeto system** command.

Detailed Steps

Command	Purpose
<pre>cluster remove unit unit_name</pre> <p>Example:</p> <pre>ciscoasa(config)# cluster remove unit ?</pre> <p>Current active units in the cluster: asa2</p> <pre>ciscoasa(config)# cluster remove unit asa2</pre> <p>WARNING: Clustering will be disabled on unit asa2. To bring it back to the cluster please logon to that unit and re-enable clustering</p>	<p>Removes the unit from the cluster. The bootstrap configuration remains intact, as well as the last configuration synced from the master unit, so you can later re-add the unit without losing your configuration. If you enter this command on a slave unit to remove the master unit, a new master unit is elected.</p> <p>To view member names, enter cluster remove unit ?, or enter the show cluster info command.</p>

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. Moreover, you cannot enable or disable clustering from a remote CLI connection.

Detailed Steps

	Command	Purpose
Step 1	<p>For a slave unit:</p> <pre>cluster group cluster_name no enable</pre> <p>Example:</p> <pre>ciscoasa(config)# cluster group cluster1 ciscoasa(cfg-cluster)# no enable</pre>	<p>Disables clustering. You cannot make configuration changes while clustering is enabled on a slave unit.</p>
Step 2	<pre>clear configure cluster</pre> <p>Example:</p> <pre>ciscoasa(config)# clear configure cluster</pre>	<p>Clears the cluster configuration. The ASA shuts down all interfaces including the management interface and cluster control link.</p>

	Command	Purpose
Step 3	no cluster interface-mode Example: <pre>ciscoasa(config)# no cluster interface-mode</pre>	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: <pre>ciscoasa(config)# copy backup_cluster.cfg running-config Source filename [backup_cluster.cfg]? Destination filename [running-config]? ciscoasa(config)#</pre>	Copies the backup configuration to the running configuration.
Step 5	write memory Example: <pre>ciscoasa(config)# write memory</pre>	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 8-54), 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 8-21 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, enter the **changeto system** command.

Detailed Steps

Command	Purpose
<code>cluster master unit unit_name</code>	Sets a new unit as the master unit. You will need to reconnect to the Main cluster IP address.
Example: <code>ciscoasa(config)# cluster master unit asa2</code>	To view member names, enter cluster master unit ? (to see all names except the current unit), or enter the show cluster info command.

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

Command	Purpose
<code>cluster exec [unit unit_name] command</code>	Sends a command to all members, or if you specify the unit name, a specific member.
Example: <code>ciscoasa# cluster exec show xlate</code>	To view member names, enter cluster exec unit ? (to see all names except the current unit), or enter the show cluster info command.

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:

```
ciscoasa# 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:

```
ciscoasa# 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

- [Monitoring Commands, page 8-59](#)
- [Related Commands, page 8-61](#)

Monitoring Commands

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

Command	Purpose
<code>show cluster info [conn-distribution packet-distribution health loadbalance trace]</code>	<p>With no keywords, the show cluster info command shows the status of all members of the cluster.</p> <p>The show cluster info conn-distribution and show cluster info packet-distribution commands show traffic distribution across all cluster units. These commands can help you to evaluate and adjust the external load balancer.</p> <p>The show cluster info trace command shows the debug information for further troubleshooting.</p> <p>The show cluster info health command shows the current health of interfaces, units, and the cluster overall.</p> <p>The show cluster info loadbalance command shows connection rebalance statistics.</p>
<code>show cluster {access-list conn cpu history interface-mode memory resource traffic xlate} [options]</code>	Displays aggregated data for the entire cluster. The <i>options</i> available depends on the data type.
<code>show cluster user-identity [options]</code>	Displays cluster-wide user identity information and statistics.
<code>show lacp cluster {system-mac system-id}</code>	Shows the cLACP system ID and priority.
<code>debug cluster [ccp datapath fsm general hc license rpc transport]</code>	Shows debug messages for clustering.
<code>debug lacp cluster [all ccp misc protocol]</code>	Shows debug messages for cLACP.
<code>show asp cluster counter</code>	This command is useful for datapath troubleshooting.

Example 8-1 show cluster info

```
ciscoasa# show cluster info
Cluster stbu: On
  This is "C" in state SLAVE
    ID       : 0
    Version  : 100.8(0.52)
    Serial No.: P3000000025
    CCL IP   : 10.0.0.3
    CCL MAC  : 000b.fcf8.c192
    Last join : 17:08:59 UTC Sep 26 2011
    Last leave: N/A
  Other members in the cluster:
    Unit "D" in state SLAVE
      ID       : 1
      Version  : 100.8(0.52)
```

```

Serial No.: P3000000001
CCL IP    : 10.0.0.4
CCL MAC   : 000b.fcf8.c162
Last join : 19:13:11 UTC Sep 23 2011
Last leave: N/A
Unit "A" in state MASTER
ID        : 2
Version   : 100.8(0.52)
Serial No.: JAB0815R0JY
CCL IP    : 10.0.0.1
CCL MAC   : 000f.f775.541e
Last join : 19:13:20 UTC Sep 23 2011
Last leave: N/A
Unit "B" in state SLAVE
ID        : 3
Version   : 100.8(0.52)
Serial No.: P3000000191
CCL IP    : 10.0.0.2
CCL MAC   : 000b.fcf8.c61e
Last join : 19:13:50 UTC Sep 23 2011
Last leave: 19:13:36 UTC Sep 23 2011

```

Example 8-2 show cluster info trace

```

ciscoasa# show cluster info trace
Feb 02 14:19:47.456 [DEBUG]Receive CCP message: CCP_MSG_LOAD_BALANCE
Feb 02 14:19:47.456 [DEBUG]Receive CCP message: CCP_MSG_LOAD_BALANCE
Feb 02 14:19:47.456 [DEBUG]Send CCP message to all: CCP_MSG_KEEPAALIVE from 80-1 at MASTER

```

Example 8-3 show cluster access-list

```

ciscoasa# show cluster access-list
hitcnt display order: cluster-wide aggregated result, unit-A, unit-B, unit-C, unit-D
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096) alert-interval
300
access-list 101; 122 elements; name hash: 0xe7d586b5
access-list 101 line 1 extended permit tcp 192.168.143.0 255.255.255.0 any eq www
(hitcnt=0, 0, 0, 0, 0) 0x207a2b7d
access-list 101 line 2 extended permit tcp any 192.168.143.0 255.255.255.0 (hitcnt=0, 0,
0, 0, 0) 0xfe4f4947
access-list 101 line 3 extended permit tcp host 192.168.1.183 host 192.168.43.238
(hitcnt=1, 0, 0, 0, 1) 0x7b521307
access-list 101 line 4 extended permit tcp host 192.168.1.116 host 192.168.43.238
(hitcnt=0, 0, 0, 0, 0) 0x5795c069
access-list 101 line 5 extended permit tcp host 192.168.1.177 host 192.168.43.238
(hitcnt=1, 0, 0, 1, 0) 0x51bde7ee
access list 101 line 6 extended permit tcp host 192.168.1.177 host 192.168.43.13
(hitcnt=0, 0, 0, 0, 0) 0x1e68697c
access-list 101 line 7 extended permit tcp host 192.168.1.177 host 192.168.43.132
(hitcnt=2, 0, 0, 1, 1) 0xc1ce5c49
access-list 101 line 8 extended permit tcp host 192.168.1.177 host 192.168.43.192
(hitcnt=3, 0, 1, 1, 1) 0xb6f59512
access-list 101 line 9 extended permit tcp host 192.168.1.177 host 192.168.43.44
(hitcnt=0, 0, 0, 0, 0) 0xdc104200
access-list 101 line 10 extended permit tcp host 192.168.1.112 host 192.168.43.44
(hitcnt=429, 109, 107, 109, 104)
0xce4f281d
access-list 101 line 11 extended permit tcp host 192.168.1.170 host 192.168.43.238
(hitcnt=3, 1, 0, 0, 2) 0x4143a818

```

```

access-list 101 line 12 extended permit tcp host 192.168.1.170 host 192.168.43.169
(hitcnt=2, 0, 1, 0, 1) 0xb18dfea4
access-list 101 line 13 extended permit tcp host 192.168.1.170 host 192.168.43.229
(hitcnt=1, 1, 0, 0, 0) 0x21557d71
access-list 101 line 14 extended permit tcp host 192.168.1.170 host 192.168.43.106
(hitcnt=0, 0, 0, 0, 0) 0x7316e016
access-list 101 line 15 extended permit tcp host 192.168.1.170 host 192.168.43.196
(hitcnt=0, 0, 0, 0, 0) 0x013fd5b8
access-list 101 line 16 extended permit tcp host 192.168.1.170 host 192.168.43.75
(hitcnt=0, 0, 0, 0, 0) 0x2c7dba0d

```

To display the aggregated count of in-use connections for all units, enter:

```

ciscoasa# show cluster conn count
Usage Summary In Cluster:*****
  200 in use (cluster-wide aggregated)
    c12(LOCAL):*****
  100 in use, 100 most used

    c11:*****
  100 in use, 100 most used

```

Related Commands

Command	Purpose
<code>show conn [detail]</code>	The show conn command shows whether a flow is a director, backup, or forwarder flow. For details about the different roles for a connection, see the “Connection Roles” section on page 8-18 . Use the cluster exec show conn command on any unit to view all connections. This command can show how traffic for a single flow arrives at different ASAs in the cluster. The throughput of the cluster is dependent on the efficiency and configuration of load balancing. This command provides an easy way to view how traffic for a connection is flowing through the cluster, and can help you understand how a load balancer might affect the performance of a flow.
<code>show route cluster</code> <code>debug route cluster</code>	Shows cluster information for routing.
<code>cluster exec capture</code>	To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the master unit using the cluster exec capture command, which is then automatically enabled on all of the slave units in the cluster. See the “Capturing Packets” section on page 43-2 .
<code>mac-address pool name start_mac_address - end_mac_address</code>	Creates a MAC address pool for an individual interface.
<code>prompt cluster-unit</code>	Sets the CLI prompt to include the cluster unit name. See the “Customizing a CLI Prompt” section on page 41-8 .

Command	Purpose
logging device-id	Each unit in the cluster generates syslog messages independently. You can use the logging device-id command to generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster. See the “Including the Device ID in Non-EMBLEM Format Syslog Messages” section on page 44-18.
show port-channel	Includes information about whether a port-channel is spanned.

Example 8-4 show conn

To troubleshoot the connection flow, first see connections on all units by entering the **cluster exec show conn** command on any unit. Look for flows that have the following flags: director (Y), backup (y), and forwarder (z). The following example shows an SSH connection from 172.18.124.187:22 to 192.168.103.131:44727 on all three ASAs; ASA 1 has the z flag showing it is a forwarder for the connection, ASA3 has the Y flag showing it is the director for the connection, and ASA2 has no special flags showing it is the owner. In the outbound direction, the packets for this connection enter the inside interface on ASA2 and exit the outside interface. In the inbound direction, the packets for this connection enter the outside interface on ASA 1 and ASA3, are forwarded over the cluster control link to ASA2, and then exit the inside interface on ASA2.

```
ciscoasa/ASA1/master# cluster exec show conn
ASA1 (LOCAL):*****
18 in use, 22 most used
Cluster stub connections: 0 in use, 5 most used
TCP outside 172.18.124.187:22 inside 192.168.103.131:44727, idle 0:00:00, bytes
37240828, flags z

ASA2:*****
12 in use, 13 most used
Cluster stub connections: 0 in use, 46 most used
TCP outside 172.18.124.187:22 inside 192.168.103.131:44727, idle 0:00:00, bytes
37240828, flags UIO

ASA3:*****
10 in use, 12 most used
Cluster stub connections: 2 in use, 29 most used
TCP outside 172.18.124.187:22 inside 192.168.103.131:44727, idle 0:00:03, bytes 0, flags
Y
```

The following is sample output for the **show conn detail** command:

```
ciscoasa/ASA2/slave# show conn detail
12 in use, 13 most used
Cluster stub connections: 0 in use, 46 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
      B - initial SYN from outside, b - TCP state-bypass or nailed,
      C - CTIQBE media, c - cluster centralized,
      D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside
      FIN,
      G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
      i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
      k - Skinny media, M - SMTP data, m - SIP media, n - GUP
      O - outbound data, P - inside back connection, p - Phone-proxy TFTP
      connection,
```

```

q - SQL*Net data, R - outside acknowledged FIN,
R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
V - VPN orphan, W - WAAS,
X - inspected by service module,
x - per session, Y - director stub flow, y - backup stub flow,
Z - Scansafe redirection, z - forwarding stub flow
ESP outside: 10.1.227.1/53744 NP Identity Ifc: 10.1.226.1/30604, , flags c, idle 0s,
uptime 1m21s, timeout 30s, bytes 7544, cluster sent/rcvd bytes 0/0, owners (0,255) Traffic
received at interface outside Locally received: 7544 (93 byte/s) Traffic received at
interface NP Identity Ifc Locally received: 0 (0 byte/s) UDP outside: 10.1.227.1/500 NP
Identity Ifc: 10.1.226.1/500, flags -c, idle 1m22s, uptime 1m22s, timeout 2m0s, bytes
1580, cluster sent/rcvd bytes 0/0, cluster sent/rcvd total bytes 0/0, owners (0,255)
Traffic received at interface outside Locally received: 864 (10 byte/s) Traffic received
at interface NP Identity Ifc Locally received: 716 (8 byte/s)

```

Configuration Examples for ASA Clustering

- [Sample ASA and Switch Configuration, page 8-63](#)
- [Firewall on a Stick, page 8-66](#)
- [Traffic Segregation, page 8-68](#)
- [Redundant Interface \(PBR or ECMP\), page 8-70](#)
- [Spanned EtherChannel With Backup Links, page 8-72](#)

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 8-63](#)
- [IOS Switch Configuration, page 8-65](#)

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
```

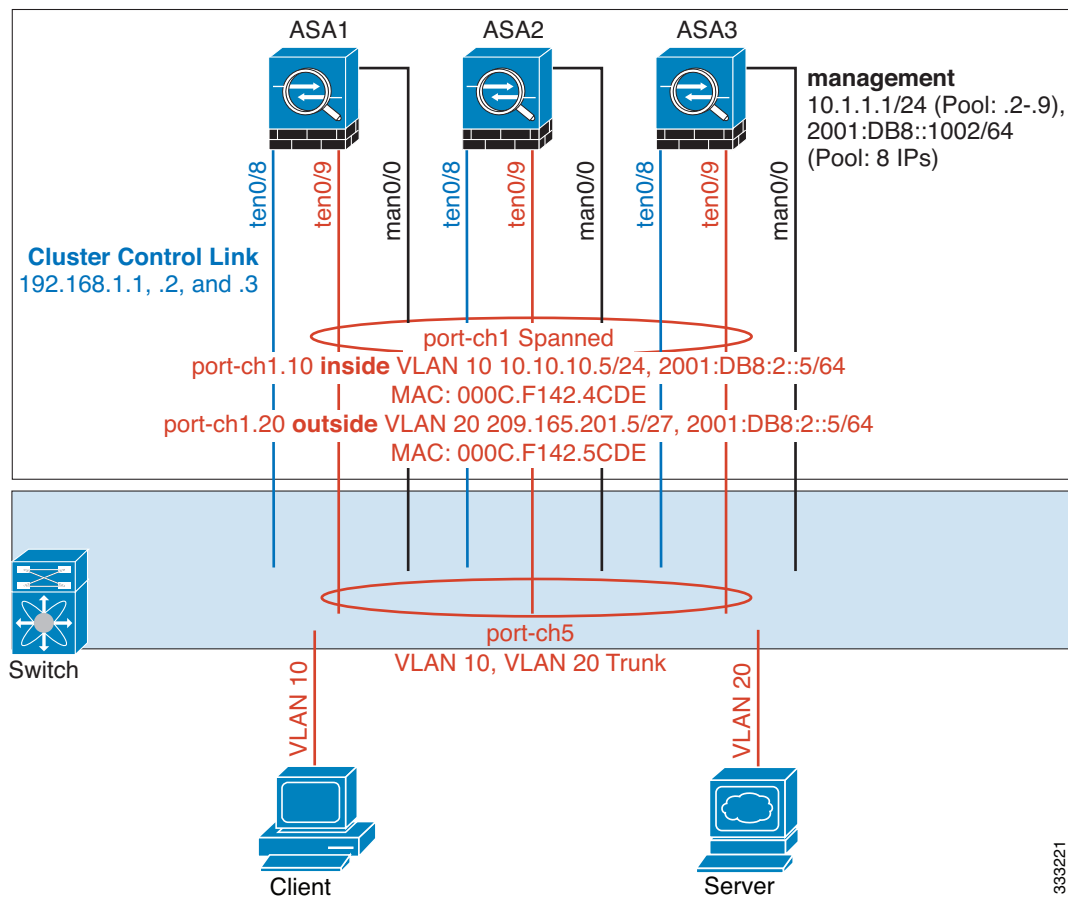
IOS Switch Configuration

```
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 asal
  cluster-interface tengigabitethernet0/8 ip 192.168.1.1 255.255.255.0
  priority 1
  key chuntheunavoidable
  enable noconfirm
```

333221

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
```

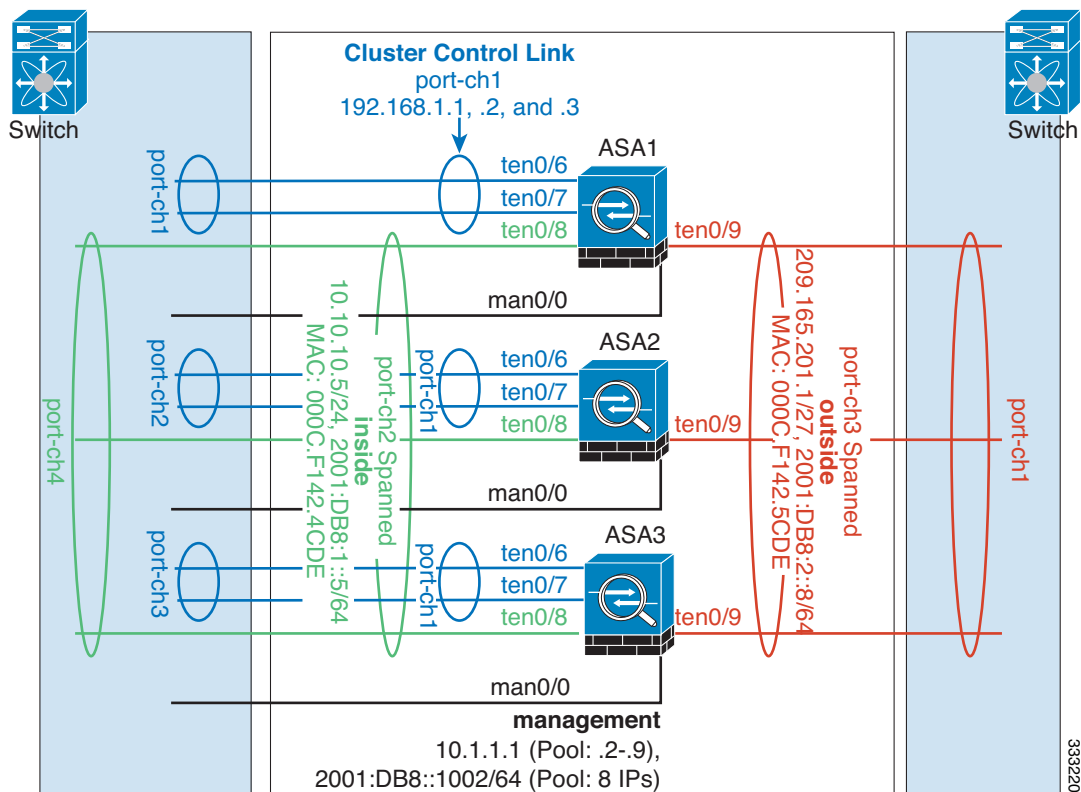
Master Interface Configuration

```
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 asal
  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

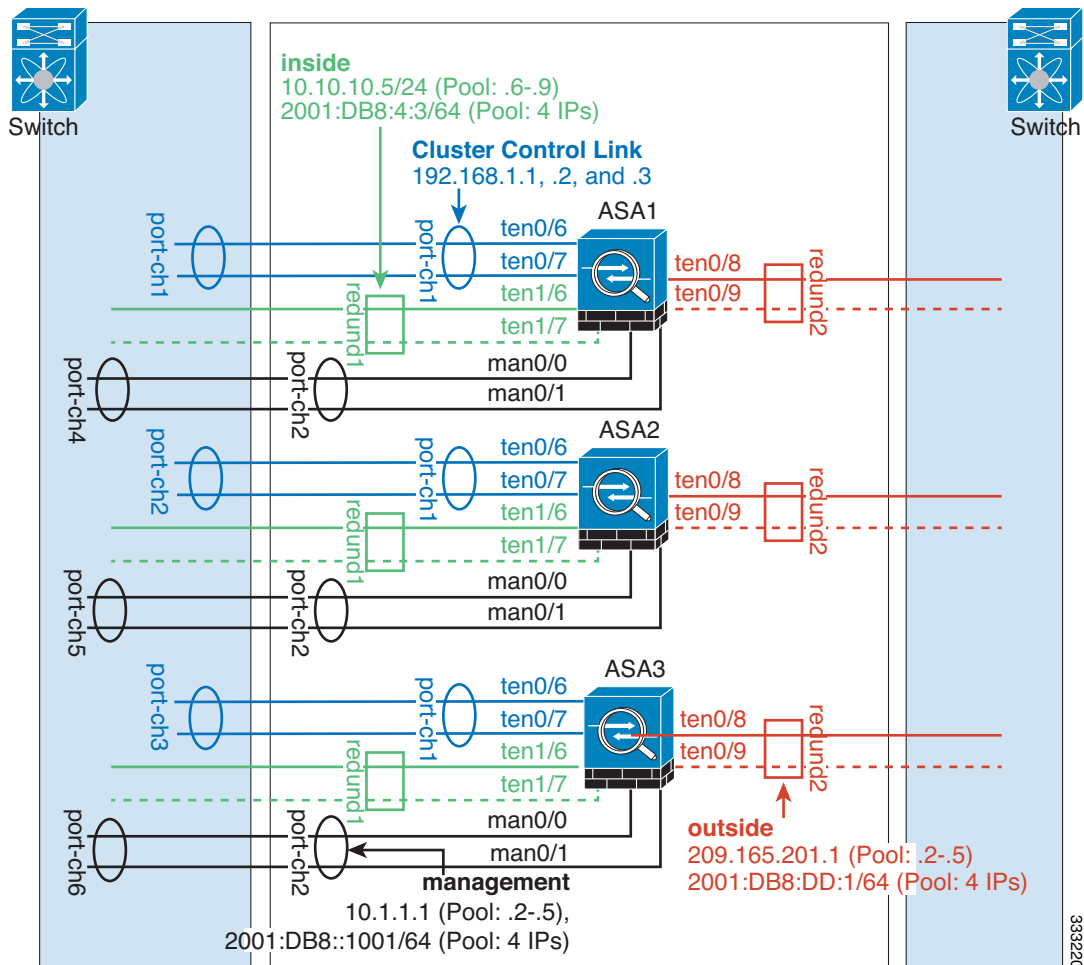
```

```

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

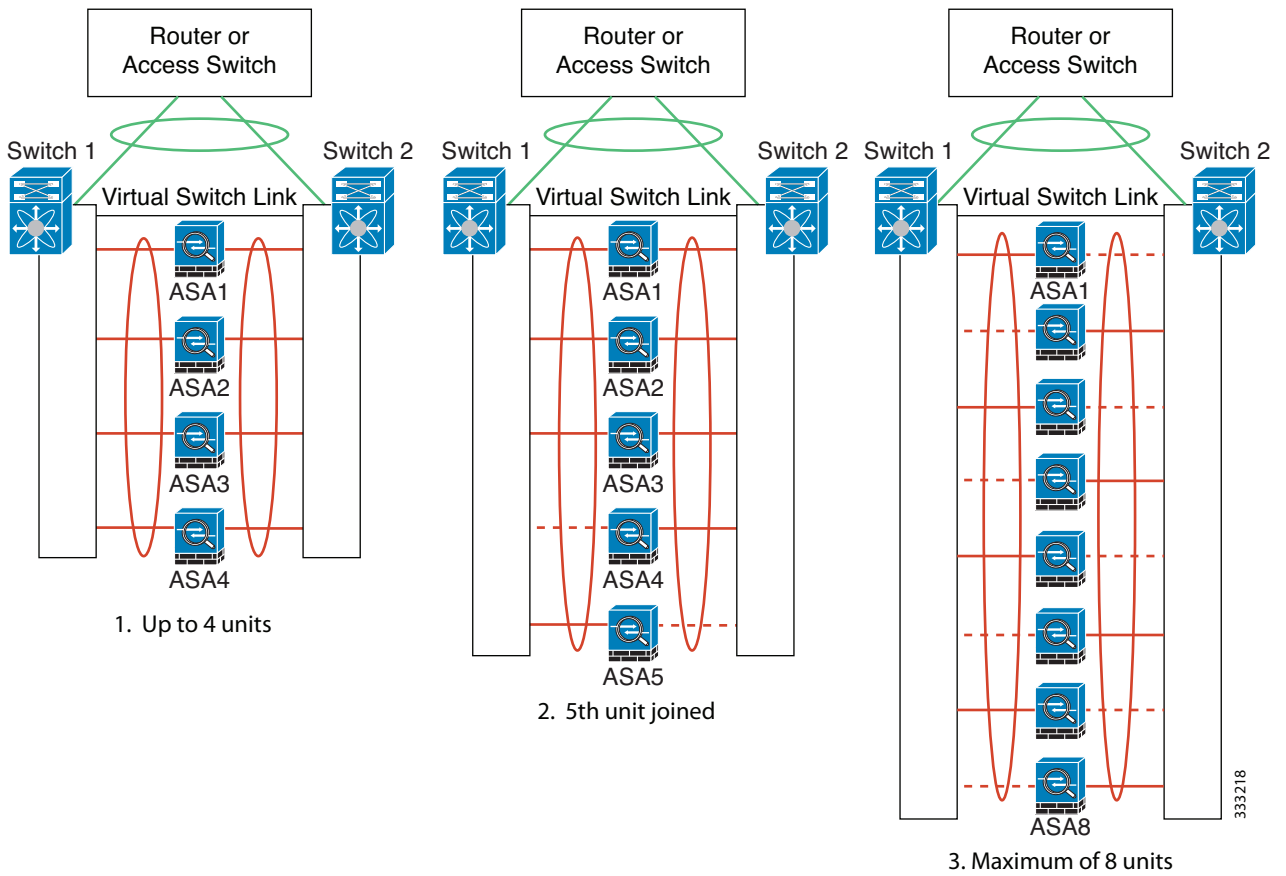
ip local pool outside 209.165.201.2-209.165.201.5
ipv6 local pool outsideipv6 2001:DB8:DD:2/64 4

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:

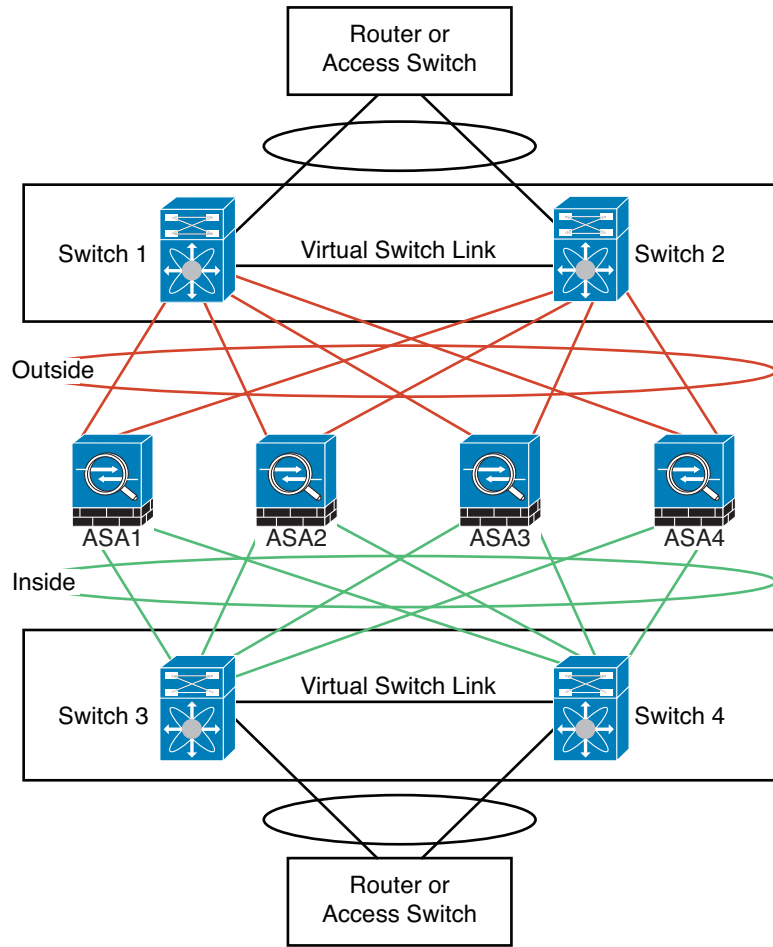


The principle is to first maximize the number of active ports in the channel, and secondly keep the number of active primary ports and the number of active secondary ports in balance. Note that when a 5th unit joins the cluster, traffic is not balanced evenly between all units.

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 8-3 lists each feature change and the platform release in which it was implemented.

Table 8-3 Feature History for Clustering

Feature Name	Platform Releases	Feature Information
ASA Clustering for the ASA 5580 and 5585-X	9.0(1)	<p>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.</p> <p>We introduced or modified the following commands: channel-group, clacp system-mac, clear cluster info, clear configure cluster, cluster exec, cluster group, cluster interface-mode, cluster-interface, conn-rebalance, console-replicate, cluster master unit, cluster remove unit, debug cluster, debug lacp cluster, enable (cluster group), health-check, ip address, ipv6 address, key (cluster group), local-unit, mac-address (interface), mac-address pool, mtu cluster, port-channel span-cluster, priority (cluster group), prompt cluster-unit, show asp cluster counter, show asp table cluster chash-table, show cluster, show cluster info, show cluster user-identity, show lacp cluster, show running-config cluster.</p>
ASA 5500-X support for clustering	9.1(4)	<p>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.</p> <p>We did not modify any ASDM screens.</p>
Improved VSS and vPC support for health check monitoring	9.1(4)	<p>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.</p> <p>We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster</p>
Support for cluster members at different geographical locations (inter-site); Individual Interface mode only	9.1(4)	<p>You can now place cluster members at different geographical locations when using individual interface mode.</p> <p>We did not modify any ASDM screens.</p>



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 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)

For ASA 5505 configuration, see [Chapter 10, “Starting Interface Configuration \(ASA 5505\).”](#)

For multiple context mode, complete all tasks in this section in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

For ASA cluster interfaces, which have special requirements, see [Chapter 8, “Configuring a Cluster of ASAs.”](#)

This chapter includes the following sections:

- [Information About Starting ASA 5510 and Higher Interface Configuration, page 9-2](#)
- [Licensing Requirements for ASA 5510 and Higher Interfaces, page 9-10](#)
- [Guidelines and Limitations, page 9-12](#)
- [Default Settings, page 9-14](#)
- [Starting Interface Configuration \(ASA 5510 and Higher\), page 9-15](#)
- [Monitoring Interfaces, page 9-36](#)
- [Configuration Examples for ASA 5510 and Higher Interfaces, page 9-36](#)
- [Where to Go Next, page 9-37](#)
- [Feature History for ASA 5510 and Higher Interfaces, page 9-38](#)

Information About Starting ASA 5510 and Higher Interface Configuration

This section includes the following topics:

- [Auto-MDI/MDIX Feature, page 9-2](#)
- [Interfaces in Transparent Mode, page 9-2](#)
- [Management Interface, page 9-2](#)
- [Redundant Interfaces, page 9-5](#)
- [EtherChannels, page 9-5](#)
- [Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size, page 9-8](#)

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 12-2.

Management Interface

- [Management Interface Overview, page 9-2](#)
- [Management Slot/Port Interface, page 9-3](#)
- [Using Any Interface for Management-Only Traffic, page 9-3](#)
- [Management Interface for Transparent Mode, page 9-4](#)
- [No Support for Redundant Management Interfaces, page 9-4](#)
- [Management 0/0 Interface on the ASA 5512-X through ASA 5555-X, page 9-4](#)

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 41, “Configuring Management Access.”](#)

Management *Slot/Port* Interface

Table 9-1 shows the Management interfaces per model.-

Table 9-1 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 (the **management-only** command). 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 3-18 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 (see the **management-only** command).

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 11-10](#) or the [“Configuring Multiple Contexts” section on page 6-15](#)). 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 9-5](#)
- [Connecting to an EtherChannel on Another Device, page 9-5](#)
- [Link Aggregation Control Protocol, page 9-6](#)
- [Load Balancing, page 9-7](#)
- [EtherChannel MAC Address, page 9-8](#)

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.

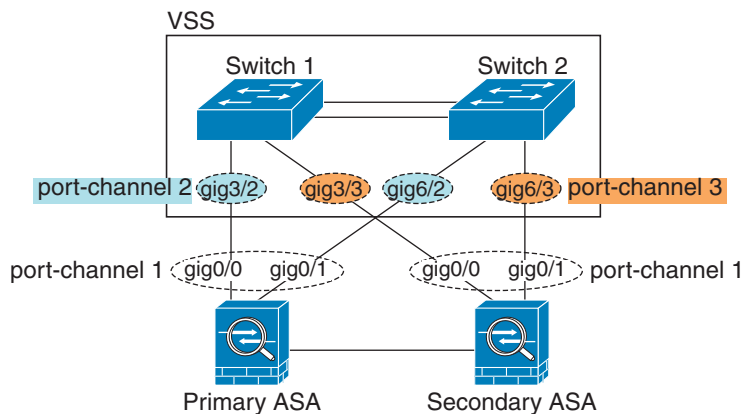
When the switch is part of a Virtual Switching System (VSS), then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch (see [Figure 9-1](#)).

Figure 9-1 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 9-1](#)). 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 9-2 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 (LACPDU) 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 9-32](#)). 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 9-2](#) shows the load balancing amounts per interface for each number of active interfaces. The active interfaces in **bold** have even distribution.

Table 9-2 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 9-8](#)
- [Default MTU, page 9-8](#)
- [Path MTU Discovery, page 9-9](#)
- [Setting the MTU and Jumbo Frames, page 9-9](#)
- [TCP Maximum Segment Size Overview, page 9-9](#)
- [Default TCP MSS, page 9-9](#)
- [Setting the TCP MSS for VPN and Non-VPN Traffic, page 9-9](#)
- [Examples, page 9-10](#)

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 9-35](#) 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 11-10. For multiple context mode, set the MTU within each context.

See the [“Enabling Jumbo Frame Support \(Supported Models\)”](#) section on page 9-35. 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 11-10. 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.

Examples

The following example enables jumbo frames, increases the MTU on all interfaces, and disables the TCP MSS for non-VPN traffic:

```
jumbo frame-reservation
mtu inside 9216
mtu outside 9216
sysopt connection tcpmss 0
```

The following example enables jumbo frames, increases the MTU on all interfaces, and changes the TCP MSS for VPN traffic:

```
jumbo frame-reservation
mtu inside 9216
mtu outside 9216
sysopt connection tcpmss 9096
```

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

Model	License Requirement
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
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

- For an interface to count against the VLAN limit, you must assign a VLAN to it. For example:

```
interface gigabitethernet 0/0.100
vlan 100
```
- The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** command defined in the configuration counts against this limit. For example, both of the following interfaces count even if the GigabitEthernet 0/0 interface is defined as part of port-channel 1:

```
interface gigabitethernet 0/0
and
interface port-channel 1
```

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

In multiple context mode, configure the physical interfaces in the system execution space according to the “[Starting Interface Configuration \(ASA 5510 and Higher\)](#)” section on page 9-15. Then, configure the logical interface parameters in the context execution space according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)

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 using the **monitor-interface** command; be sure to reference the logical redundant interface name. 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 8-38](#).
- To configure an individual cluster interface, see the [“Configuring Individual Interfaces \(Recommended for the Management Interface\)” section on page 8-36](#).

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 9-13](#).
- For clustering guidelines, see the [“Clustering Guidelines” section on page 9-13](#).

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 LACPDU s 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 LACPDU s. 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 9-13](#).
- For clustering guidelines, see the [“Clustering Guidelines” section on page 9-13](#).

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 3-18](#).

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 9-15](#)
- [Converting In-Use Interfaces to a Redundant or EtherChannel Interface, page 9-16](#)
- [Enabling the Physical Interface and Configuring Ethernet Parameters, page 9-26](#)
- [Configuring a Redundant Interface, page 9-28](#)
- [Configuring an EtherChannel, page 9-30](#)
- [Configuring VLAN Subinterfaces and 802.1Q Trunking, page 9-33](#)
- [Enabling Jumbo Frame Support \(Supported Models\), page 9-35](#)

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 to minimize disruption. See the [“Converting In-Use Interfaces to a Redundant or EtherChannel Interface”](#) section on page 9-16.

To start configuring interfaces, perform the following steps:

-
- Step 1** (Multiple context mode) Complete all tasks in this section in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.
- Step 2** Enable the physical interface, and optionally change Ethernet parameters. See the [“Enabling the Physical Interface and Configuring Ethernet Parameters”](#) section on page 9-26.

Physical interfaces are disabled by default.

- Step 3** (Optional) Configure redundant interface pairs. See the [“Configuring a Redundant Interface”](#) section on page 9-28.

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 9-30. 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 9-33.

- 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 9-35.

- 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 6, [“Configuring Multiple Context Mode”](#):
- To assign interfaces to contexts, see the [“Configuring a Security Context”](#) section on page 6-19.
 - (Optional) To automatically assign unique MAC addresses to context interfaces, see the [“Automatically Assigning MAC Addresses to Context Interfaces”](#) section on page 6-24.

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 11-10.

- Step 8** Complete the interface configuration according to Chapter 11, [“Completing Interface Configuration \(Routed Mode\)”](#), or Chapter 12, [“Completing Interface Configuration \(Transparent Mode\)”](#).

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 9-28 and the [“Configuring an EtherChannel”](#) section on page 9-30 for more information.

- [Detailed Steps \(Single Mode\)](#), page 9-16
- [Detailed Steps \(Multiple Mode\)](#), page 9-22

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 entering the **no failover** command.
- Step 3** Copy the running configuration by entering the **more system:running-config** command 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.
- 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 9-30](#).

Step 8 At the ASA CLI prompt, perform the following steps depending on your connection (console or remote).

- Console connection:
 - a. Copy the entire new configuration to the clipboard, including the altered interface section.
 - b. Clear the running configuration by entering:


```
ciscoasa(config)# clear configure all
```

Traffic through the ASA stops at this point.
 - c. Paste in the new configuration at the prompt.

Traffic through the ASA resumes.
- Remote connection:
 - a. Save the new configuration to a TFTP or FTP server, so you can copy it to the startup configuration on the ASA. For example, you can run a TFTP or FTP server on your PC.
 - b. Clear the startup configuration by entering:


```
ciscoasa(config)# write erase
```
 - c. Copy the new configuration to the startup configuration by entering:


```
ciscoasa(config)# copy url startup-config
```

See the [“Copying a File to the ASA” section on page 42-17](#)
 - d. Reload the ASA using the **reload** command. Do not save the running configuration.

Step 9 Reenable failover by entering the **failover** command.

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 entering the **no failover** command.
- Step 3** In the system, copy the running configuration by entering the **more system:running-config** command 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 Configurations or Other Files”](#) section on page 42-25.

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 9-28 or the [“Configuring an EtherChannel”](#) section on page 9-30. 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 6-19.

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, if your contexts are on an FTP server, copy over the existing files (making backups as desired) using FTP. If your contexts are in flash memory, you can use the **copy** command and run a TFTP or FTP server on your PC, or use secure copy. See the [“Copying a File to the ASA” section on page 42-17](#). This change only affects the startup configuration; the running configuration is still using the old context configuration.
- Step 9** At the ASA system CLI prompt, perform the following steps depending on your connection (console or remote).
- Console connection:
 - a. Copy the entire new system configuration to the clipboard, including the altered interface section.
 - b. Clear the running configuration (both system and contexts) by entering:


```
ciscoasa(config)# clear configure all
```

Traffic through the ASA stops at this point.
 - c. Paste in the new system configuration at the prompt.

All of the new context configurations now reload. When they are finished reloading, traffic through the ASA resumes.
 - Remote connection:
 - a. Save the new system configuration to a TFTP or FTP server, so you can copy it to the startup configuration on the ASA. For example, you can run a TFTP or FTP server on your PC.
 - b. Clear the startup configuration by entering:


```
ciscoasa(config)# write erase
```
 - c. Copy the new system configuration to the startup configuration by entering:


```
ciscoasa(config)# copy url startup-config
```

See the [“Copying a File to the ASA” section on page 42-17](#)
 - d. Reload the ASA using the **reload** command. Do not save the running configuration.
- Step 10** Reenable failover by entering the **failover** command.

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. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	<p>interface <i>physical_interface</i></p> <p>Example: <pre>ciscoasa(config)# interface gigabitethernet 0/0</pre></p>	<p>Specifies the interface you want to configure.</p> <p>where the <i>physical_interface</i> ID includes the type, slot, and port number as <i>type[slot/port]</i>.</p> <p>The physical interface types include the following:</p> <ul style="list-style-type: none"> • ethernet • gigabitethernet • tengigabitethernet • management <p>Enter the type followed by <i>slot/port</i>, for example, gigabitethernet0/1 or ethernet 0/1. A space is optional between the type and the slot/port.</p>
Step 2	<p>(Optional)</p> <p>media-type sfp</p> <p>Example: <pre>ciscoasa(config-if)# media-type sfp</pre></p>	<p>Sets the media type to SFP, if available for your model. To restore the default RJ-45, enter the media-type rj45 command.</p>
Step 3	<p>(Optional)</p> <p>speed {auto 10 100 1000 nonegotiate}</p> <p>Example: <pre>ciscoasa(config-if)# speed 100</pre></p>	<p>Sets the speed.</p> <p>For copper interfaces, the default setting is auto.</p> <p>For SFP interfaces, the default setting is no speed nonegotiate, which sets the speed to the maximum speed and enables link negotiation for flow-control parameters and remote fault information. The nonegotiate keyword is the only keyword available for SFP interfaces. The speed nonegotiate command disables link negotiation.</p>

	Command	Purpose
Step 4	(Optional) <pre>duplex {auto full half}</pre> Example: <pre>ciscoasa(config-if)# duplex full</pre>	Sets the duplex for copper interfaces. The auto setting is the default. Note The duplex setting for an EtherChannel interface must be Full or Auto.
Step 5	(Optional) <pre>flowcontrol send on [low_water high_water pause_time] [noconfirm]</pre> Example: <pre>ciscoasa(config-if)# flowcontrol send on 95 200 10000</pre>	Enables pause (XOFF) frames for flow control on 1-Gigabit and 10-Gigabit Ethernet interfaces. 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 <i>high_water</i> 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 <i>low_water</i> 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 <i>pause_time</i> 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. When you use this command, you see the following warning: <pre>Changing flow-control parameters will reset the interface. Packets may be lost during the reset. Proceed with flow-control changes?</pre> To change the parameters without being prompted, use the noconfirm keyword. Note Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.
Step 6	<pre>no shutdown</pre> Example: <pre>ciscoasa(config-if)# no shutdown</pre>	Enables the interface. To disable the interface, enter the shutdown command. If you enter the shutdown command, you also shut down all subinterfaces. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it.

What to Do Next

Optional Tasks:

- Configure redundant interface pairs. See the [“Configuring a Redundant Interface”](#) section on page 9-28.
- Configure an EtherChannel. See the [“Configuring an EtherChannel”](#) section on page 9-30.

- Configure VLAN subinterfaces. See the “[Configuring VLAN Subinterfaces and 802.1Q Trunking](#)” section on page 9-33.

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 6-15.
- For single context mode, complete the interface configuration. See Chapter 11, “[Completing Interface Configuration \(Routed Mode\)](#),” or Chapter 12, “[Completing Interface Configuration \(Transparent Mode\)](#).”

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 9-28
- [Changing the Active Interface](#), page 9-30

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 9-13.

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 using the **no nameif** command.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.



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

	Command	Purpose
Step 1	interface redundant <i>number</i> Example: <pre>ciscoasa(config)# interface redundant 1</pre>	Adds the logical redundant interface, where the <i>number</i> argument is an integer between 1 and 8. Note You need to add at least one member interface to the redundant interface before you can configure logical parameters for it such as a name.
Step 2	member-interface <i>physical_interface</i> Example: <pre>ciscoasa(config-if)# member-interface management 0/0</pre>	Adds the first member interface to the redundant interface. See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section on page 9-26 for a description of the physical interface ID. Redundant interfaces do not support Management <i>slot/port</i> interfaces as members. After you add the interface, any configuration for it (such as an IP address) is removed.
Step 3	member-interface <i>physical_interface</i> Example: <pre>ciscoasa(config-if)# member-interface management 1/0</pre>	Adds the second member interface to the redundant interface. Make sure the second interface is the same physical type as the first interface. To remove a member interface, enter the no member-interface <i>physical_interface</i> command. You cannot remove both member interfaces from the redundant interface; the redundant interface requires at least one member interface.

Examples

The following example creates two redundant interfaces:

```
ciscoasa(config)# interface redundant 1
ciscoasa(config-if)# member-interface gigabitethernet 0/0
ciscoasa(config-if)# member-interface gigabitethernet 0/1
ciscoasa(config-if)# interface redundant 2
ciscoasa(config-if)# member-interface gigabitethernet 0/2
ciscoasa(config-if)# member-interface gigabitethernet 0/3
```

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See the [“Configuring VLAN Subinterfaces and 802.1Q Trunking”](#) section on page 9-33.

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 6-15.
- For single context mode, complete the interface configuration. See the [Chapter 11, “Completing Interface Configuration \(Routed Mode\)”](#), or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\)”](#).

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:

```
ciscoasa# show interface redundantnumber detail | grep Member
```

For example:

```
ciscoasa# show interface redundant1 detail | grep Member
Members GigabitEthernet0/3 (Active), GigabitEthernet0/2
```

To change the active interface, enter the following command:

```
ciscoasa# redundant-interface redundantnumber active-member physical_interface
```

where the **redundantnumber** 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 9-30](#)
- [Customizing the EtherChannel, page 9-32](#)

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 8-38](#) instead of this procedure.
- See also the [“EtherChannel Guidelines” section on page 9-13](#).

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 using the **no nameif** command.

- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

**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

	Command	Purpose
Step 1	<p>interface <i>physical_interface</i></p> <p>Example: ciscoasa(config)# interface gigabitethernet 0/0</p>	<p>Specifies the interface you want to add to the channel group, where the <i>physical_interface</i> ID includes the type, slot, and port number as <i>type[slot/port</i>. This first interface in the channel group determines the type and speed for all other interfaces in the group.</p> <p>In transparent mode, if you create a channel group with multiple Management interfaces, then you can use this EtherChannel as the management-only interface.</p>
Step 2	<p>channel-group <i>channel_id</i> mode {active passive on}</p> <p>Example: ciscoasa(config-if)# channel-group 1 mode active</p>	<p>Assigns this physical interface to an EtherChannel with the <i>channel_id</i> between 1 and 48. If the port-channel interface for this channel ID does not yet exist in the configuration, one will be added:</p> <p>interface port-channel <i>channel_id</i></p> <p>We recommend using active mode. For information about active, passive, and on modes, see the “Link Aggregation Control Protocol” section on page 9-6.</p>
Step 3	<p>(Optional)</p> <p>lacp port-priority <i>number</i></p> <p>Example: ciscoasa(config-if)# lacp port-priority 12345</p>	<p>Sets the priority for a physical interface in the channel group 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.</p> <p>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 lacp port-priority value be 12345 on the 1/3 interface vs. the default 32768 on the 0/7 interface.</p> <p>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 the lacp system-priority command in the “Customizing the EtherChannel” section on page 9-32.</p>
Step 4	<p>Repeat steps 1 through 3 for each interface you want to add to the channel group.</p>	<p>Each interface in the channel group must be the same type and speed. Half duplex is not supported. If you add an interface that does not match, it will be placed in a suspended state.</p>

What to Do Next

Optional Tasks:

- Customize the EtherChannel interface. See the “Customizing the EtherChannel” section on page 9-32.
- Configure VLAN subinterfaces. See the “Configuring VLAN Subinterfaces and 802.1Q Trunking” section on page 9-33.

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 6-15.
- For single context mode, complete the interface configuration. See the Chapter 11, “Completing Interface Configuration (Routed Mode),” or Chapter 12, “Completing Interface Configuration (Transparent Mode).”

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

	Command	Purpose
Step 1	<p>interface port-channel <i>channel_id</i></p> <p>Example: ciscoasa(config)# interface port-channel 1</p>	<p>Specifies the port-channel interface. This interface was created automatically when you added an interface to the channel group. If you have not yet added an interface, then this command creates the port-channel interface.</p> <p>Note You need to add at least one member interface to the port-channel interface before you can configure logical parameters for it such as a name.</p>
Step 2	<p>lACP max-bundle <i>number</i></p> <p>Example: ciscoasa(config-if)# lACP max-bundle 6</p>	<p>Specifies the maximum number of active interfaces allowed in the channel group, between 1 and 8. The default is 8.</p>
Step 3	<p>port-channel min-bundle <i>number</i></p> <p>Example: ciscoasa(config-if)# port-channel min-bundle 2</p>	<p>Specifies the minimum number of active interfaces required for the port-channel interface to become active, between 1 and 8. The default is 1. If the active interfaces in the channel group falls below this value, then the port-channel interface goes down, and could trigger a device-level failover.</p>

	Command	Purpose
Step 4	<pre>port-channel load-balance {dst-ip dst-ip-port dst-mac dst-port src-dst-ip src-dst-ip-port src-dst-mac src-dst-port src-ip src-ip-port src-mac src-port vlan-dst-ip vlan-dst-ip-port vlan-only vlan-src-dst-ip vlan-src-dst-ip-port vlan-src-ip vlan-src-ip-port}</pre> <p>Example: ciscoasa(config-if)# port-channel load-balance src-dst-mac</p>	<p>Configures the load-balancing algorithm. By default, the ASA balances the packet load on interfaces according to the source and destination IP address (src-dst-ip) of the packet. If you want to change the properties on which the packet is categorized, use this command. 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 9-7.</p>
Step 5	<pre>lACP system-priority number</pre> <p>Example: ciscoasa(config)# lACP system-priority 12345</p>	<p>Sets the LACP system priority, from 1 to 65535. The default is 32768. The higher the number, the lower the priority. This command is global for the ASA.</p> <p>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. For interface priorities within an EtherChannel, see the lACP port-priority command in the “Adding Interfaces to the EtherChannel” section on page 9-30.</p>
Step 6	<p>(Optional)</p> <p>You can set the Ethernet properties for the port-channel interface to override the properties set on the individual interfaces.</p>	<p>This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group. See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section on page 9-26 for Ethernet commands.</p>

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See the [“Configuring VLAN Subinterfaces and 802.1Q Trunking”](#) section on page 9-33.

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 6-15.
- For single context mode, complete the interface configuration. See the [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)

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 9-10.
- 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 leaving out the **nameif** command. If you want to let the physical or redundant interface pass untagged packets, you can configure the **nameif** command as usual. See [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\),”](#) 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. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	<p>interface {<i>physical_interface</i> redundant number port-channel number}.<i>subinterface</i></p> <p>Example: <pre>ciscoasa(config)# interface gigabitethernet 0/1.100</pre></p>	<p>Specifies the new subinterface. See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>The <i>subinterface</i> ID is an integer between 1 and 4294967293.</p>
Step 2	<p>vlan <i>vlan_id</i></p> <p>Example: <pre>ciscoasa(config-subif)# vlan 101</pre></p>	<p>Specifies the VLAN for the subinterface. The <i>vlan_id</i> is an integer between 1 and 4094. Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information.</p> <p>You can only assign a single VLAN to a subinterface, and you cannot assign the same VLAN to multiple subinterfaces. You cannot assign a VLAN to the physical interface. Each subinterface must have a VLAN ID before it can pass traffic. To change a VLAN ID, you do not need to remove the old VLAN ID with the no option; you can enter the vlan command with a different VLAN ID, and the ASA changes the old ID.</p>

What to Do Next

(Optional) Enable jumbo frame support according to the [“Enabling Jumbo Frame Support \(Supported Models\)”](#) section on page 9-35.

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 9-8 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 using the **mtu** command. See the [“Configuring the MAC Address, MTU, and TCP MSS”](#) section on page 11-10. 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 (**sysopt connection tcpmss 0**), or to increase it in accord with the MTU according to the [“Configuring the MAC Address, MTU, and TCP MSS”](#) section on page 11-10.

Detailed Steps

Command	Purpose
jumbo-frame reservation Example: ciscoasa(config)# jumbo-frame reservation	Enables jumbo frame support. To disable jumbo frames, use the no form of this command.

Examples

The following example enables jumbo frame reservation, saves the configuration, and reloads the ASA:

```
ciscoasa(config)# jumbo-frame reservation
WARNING: this command will take effect after the running-config is saved
and the system has been rebooted. Command accepted.

ciscoasa(config)# write memory
Building configuration...
Cryptochecksum: 718e3706 4edb11ea 69af58d0 0a6b7cb5

70291 bytes copied in 3.710 secs (23430 bytes/sec)
[OK]
ciscoasa(config)# reload
Proceed with reload? [confirm] Y
```

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
<code>show interface</code>	Displays interface statistics.
<code>show interface ip brief</code>	Displays interface IP addresses and status.
<code>show lacp</code> <i>[[channel_group_number] {counters internal neighbor} sys-id]</i>	For EtherChannel, displays LACP information such as traffic statistics, system identifier and neighbor details.
<code>show port-channel</code> <i>[channel_group_number] [brief detail port protocol summary]</i>	For EtherChannel, displays EtherChannel information in a detailed and one-line summary form. This command also displays the port and port-channel information.
<code>show port-channel</code> <i>channel_group_number load-balance [hash-result {ip ipv6 l4port mac mixed vlan-only} parameters]</i>	For EtherChannel, displays port-channel load-balance information along with the hash result and member interface selected for a given set of parameters.

Configuration Examples for ASA 5510 and Higher Interfaces

This section includes the following topics:

- [Physical Interface Parameters Example, page 9-36](#)
- [Subinterface Parameters Example, page 9-37](#)
- [Multiple Context Mode Example, page 9-37](#)
- [EtherChannel Example, page 9-37](#)

Physical Interface Parameters Example

The following example configures parameters for the physical interface in single mode:

```
interface gigabitethernet 0/1
```

```
speed 1000
duplex full
no shutdown
```

Subinterface Parameters Example

The following example configures parameters for a subinterface in single mode:

```
interface gigabitEthernet 0/1.1
vlan 101
no shutdown
```

Multiple Context Mode Example

The following example configures interface parameters in multiple context mode for the system configuration, and allocates the gigabitEthernet 0/1.1 subinterface to contextA:

```
interface gigabitEthernet 0/1
speed 1000
duplex full
no shutdown
interface gigabitEthernet 0/1.1
vlan 101
context contextA
allocate-interface gigabitEthernet 0/1.1
```

EtherChannel Example

The following example configures three interfaces as part of an EtherChannel. It also sets the system priority to be a higher priority, and GigabitEthernet 0/2 to be a higher priority than the other interfaces in case more than eight interfaces are assigned to the EtherChannel.

```
lacp system-priority 1234
interface GigabitEthernet0/0
channel-group 1 mode active
interface GigabitEthernet0/1
channel-group 1 mode active
interface GigabitEthernet0/2
lacp port-priority 1234
channel-group 1 mode passive
interface Port-channel1
lacp max-bundle 4
port-channel min-bundle 2
port-channel load-balance dst-ip
```

Where to Go Next

- For multiple context mode:
 - a. Assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See [Chapter 6, “Configuring Multiple Context Mode.”](#)

- b. Complete the interface configuration according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)
- For single context mode, complete the interface configuration according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\),”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\).”](#)

Feature History for ASA 5510 and Higher Interfaces

Table 9-3 lists the release history for this feature.

Table 9-3 Feature History for Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: <ul style="list-style-type: none"> • 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. Use the speed command to change the speed on the interface and use the show interface command to see what speed is currently configured for each interface.

Table 9-3 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
Jumbo packet support for the ASA 5580	8.1(1)	<p>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.</p> <p>This feature is also supported on the ASA 5585-X.</p> <p>We introduced the following command: jumbo-frame reservation.</p>
Increased VLANs for the ASA 5580	8.1(2)	<p>The number of VLANs supported on the ASA 5580 are increased from 100 to 250.</p>
Support for Pause Frames for Flow Control on the ASA 5580 10-Gigabit Ethernet Interfaces	8.2(2)	<p>You can now enable pause (XOFF) frames for flow control.</p> <p>This feature is also supported on the ASA 5585-X.</p> <p>We introduced the following command: flowcontrol.</p>
Support for Pause Frames for Flow Control on 1-Gigabit Ethernet Interfaces	8.2(5)/8.4(2)	<p>You can now enable pause (XOFF) frames for flow control for 1-Gigabit interfaces on all models.</p> <p>We modified the following command: flowcontrol.</p>
EtherChannel support	8.4(1)	<p>You can configure up to 48 802.3ad EtherChannels of eight active interfaces each.</p> <p>We introduced the following commands: channel-group, lACP port-priority, interface port-channel, lACP max-bundle, port-channel min-bundle, port-channel load-balance, lACP system-priority, clear lACP counters, show lACP, show port-channel.</p> <p>Note EtherChannel is not supported on the ASA 5505.</p>



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 10-13](#).

This chapter includes the following sections:

- [Information About ASA 5505 Interfaces](#), page 10-1
- [Licensing Requirements for ASA 5505 Interfaces](#), page 10-4
- [Guidelines and Limitations](#), page 10-5
- [Default Settings](#), page 10-5
- [Starting ASA 5505 Interface Configuration](#), page 10-6
- [Monitoring Interfaces](#), page 10-11
- [Configuration Examples for ASA 5505 Interfaces](#), page 10-11
- [Where to Go Next](#), page 10-13
- [Feature History for ASA 5505 Interfaces](#), page 10-13

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 10-2
- [Maximum Active VLAN Interfaces for Your License](#), page 10-2
- [VLAN MAC Addresses](#), page 10-4
- [Power over Ethernet](#), page 10-4
- [Monitoring Traffic Using SPAN](#), page 10-4
- [Auto-MDI/MDIX Feature](#), page 10-4

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 10-4 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 10-1](#) for more information.
- Security Plus license—20 active VLANs.

In transparent firewall mode, you can configure the following VLANs depending on your license:

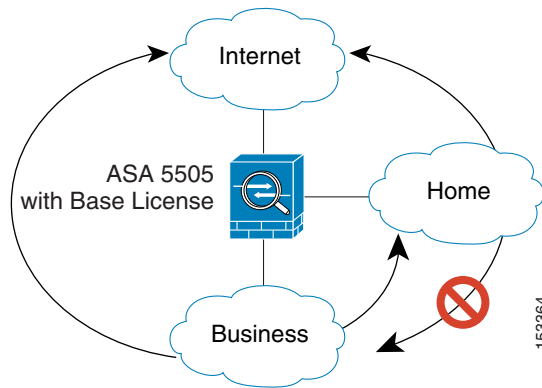
- Base license—2 active VLANs in 1 bridge group.
- Security Plus license—3 active VLANs: 2 active VLANs in 1 bridge group, and 1 active VLAN for the failover link.

**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 10-1](#) for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 10-1 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 10-2](#) for an example network.

Figure 10-2 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 11-10](#).
- 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 12-13](#).

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 using the **shutdown** command, you disable power to the device. Power is restored when you enable the port using the **no shutdown** command. See the [“Configuring and Enabling Switch Ports as Access Ports” section on page 10-7](#) for more information about shutting down a switch port.

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

Monitoring Traffic Using SPAN

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

See the **switchport monitor** command in the 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	<p>VLANs:</p> <p>Routed Mode:</p> <p>Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone)</p> <p>Security Plus License: 20</p> <p>Transparent Mode:</p> <p>Base License: 2 active VLANs in 1 bridge group.</p> <p>Security Plus License: 3 active VLANs: 2 active VLANs in 1 bridge group, and 1 active VLAN for the failover link.</p> <p>VLAN Trunks:</p> <p>Base License: None.</p> <p>Security Plus License: 8.</p>

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 3-18](#).

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 10-6](#)
- [Configuring VLAN Interfaces, page 10-6](#)
- [Configuring and Enabling Switch Ports as Access Ports, page 10-7](#)
- [Configuring and Enabling Switch Ports as Trunk Ports, page 10-9](#)

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 10-6. |
| Step 2 | Configure and enable switch ports as access ports. See the “Configuring and Enabling Switch Ports as Access Ports” section on page 10-7. |
| 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 10-9. |
| Step 4 | Complete the interface configuration according to Chapter 11, “Completing Interface Configuration (Routed Mode),” or Chapter 12, “Completing Interface Configuration (Transparent Mode).” |
-

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 10-1.

Guidelines

We suggest that you finalize your interface configuration before you enable Easy VPN.

Detailed Steps

	Command	Purpose
Step 1	<p><code>interface vlan number</code></p> <p>Example: <pre>ciscoasa(config)# interface vlan 100</pre></p>	<p>Adds a VLAN interface, where the <i>number</i> is between 1 and 4090.</p> <p>To remove this VLAN interface and all associated configuration, enter the no interface vlan command. Because this interface also includes the interface name configuration, and the name is used in other commands, those commands are also removed.</p>
Step 2	<p>(Optional for the Base license)</p> <p><code>no forward interface vlan number</code></p> <p>Example: <pre>ciscoasa(config-if)# no forward interface vlan 101</pre></p>	<p>Allows this interface to be the third VLAN by limiting it from initiating contact to one other VLAN.</p> <p>The <i>number</i> specifies the VLAN ID to which this VLAN interface cannot initiate traffic.</p> <p>With the Base license, you can only configure a third VLAN if you use this command to limit it.</p> <p>For example, you have one VLAN assigned to the outside for Internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can use the no forward interface command on the home VLAN; the business network can access the home network, but the home network cannot access the business network.</p> <p>If you already have two VLAN interfaces configured with a nameif command, be sure to enter the no forward interface command before the nameif command on the third interface; the ASA does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505.</p> <p>Note If you upgrade to the Security Plus license, you can remove this command and achieve full functionality for this interface. If you leave this command in place, this interface continues to be limited even after upgrading.</p>

What to Do Next

Configure the switch ports. See the [“Configuring and Enabling Switch Ports as Access Ports”](#) section on page 10-7 and the [“Configuring and Enabling Switch Ports as Trunk Ports”](#) section on page 10-9.

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 10-9. If you have a factory default configuration, see the [“ASA 5505 Default Configuration”](#) section on page 3-20 to check if you want to change the default interface settings according to this procedure.

For more information about ASA 5505 interfaces, see the [“Information About ASA 5505 Interfaces”](#) section on page 10-1.

**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

	Command	Purpose
Step 1	<code>interface ethernet0/port</code> Example: <code>ciscoasa(config)# interface ethernet0/1</code>	Specifies the switch port you want to configure, where <i>port</i> is 0 through 7.
Step 2	<code>switchport access vlan number</code> Example: <code>ciscoasa(config-if)# switchport access vlan 100</code>	Assigns this switch port to a VLAN, where <i>number</i> is the VLAN ID, between 1 and 4090. See the “Configuring VLAN Interfaces” section on page 10-6 to configure the VLAN interface that you want to assign to this switch port. To view configured VLANs, enter the show interface command. Note You might assign multiple switch ports to the primary or backup VLANs if the Internet access device includes Layer 2 redundancy.
Step 3	(Optional) <code>switchport protected</code> Example: <code>ciscoasa(config-if)# switchport protected</code>	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 switchport protected command 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 4	(Optional) <code>speed {auto 10 100}</code> Example: <code>ciscoasa(config-if)# speed 100</code>	Sets the speed. The auto setting is the default. If you set the speed to anything other than auto on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.
Step 5	(Optional) <code>duplex {auto full half}</code> Example: <code>ciscoasa(config-if)# duplex full</code>	Sets the duplex. 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 6	<code>no shutdown</code> Example: <code>ciscoasa(config-if)# no shutdown</code>	Enables the switch port. To disable the switch port, enter the shutdown command.

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 10-9.
- To complete the interface configuration, see [Chapter 11, “Completing Interface Configuration \(Routed Mode\)”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\)”](#).

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 10-7.

Guidelines

This switch port cannot pass traffic until you assign at least one VLAN to it, native or non-native.

Detailed Steps

	Command	Purpose
Step 1	interface ethernet0/port Example: ciscoasa(config)# interface ethernet0/1	Specifies the switch port you want to configure, where <i>port</i> is 0 through 7.
Step 2	To assign VLANs to this trunk, do one or more of the following: switchport trunk allowed vlan vlan_range Example: ciscoasa(config)# switchport trunk allowed vlan 100-200	Identifies one or more VLANs that you can assign to the trunk port, where the <i>vlan_range</i> (with VLANs between 1 and 4090) can be identified in one of the following ways: <ul style="list-style-type: none"> • A single number (n) • A range (n-x) • Separate numbers and ranges by commas, for example: 5,7-10,13,45-100 You can enter spaces instead of commas, but the command is saved to the configuration with commas. You can include the native VLAN in this command, but it is not required; the native VLAN is passed whether it is included in this command or not.

Command	Purpose
<pre>switchport trunk native vlan <i>vlan_id</i></pre> <p>Example: <pre>ciscoasa(config-if)# switchport trunk native vlan 100</pre></p>	<p>Assigns a native VLAN to the trunk, where the <i>vlan_id</i> is a single VLAN ID between 1 and 4090.</p> <p>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.</p> <p>Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.</p>
<p>Step 3</p> <pre>switchport mode trunk</pre> <p>Example: <pre>ciscoasa(config-if)# switchport mode trunk</pre></p>	<p>Makes this switch port a trunk port. To restore this port to access mode, enter the switchport mode access command.</p>
<p>Step 4</p> <p>(Optional)</p> <pre>switchport protected</pre> <p>Example: <pre>ciscoasa(config-if)# switchport protected</pre></p>	<p>Prevents the switch port from communicating with other protected switch ports on the same VLAN.</p> <p>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 switchport protected command 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.</p>
<p>Step 5</p> <p>(Optional)</p> <pre>speed {auto 10 100}</pre> <p>Example: <pre>ciscoasa(config-if)# speed 100</pre></p>	<p>Sets the speed. The auto setting is the default. If you set the speed to anything other than auto on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.</p>
<p>Step 6</p> <p>(Optional)</p> <pre>duplex {auto full half}</pre> <p>Example: <pre>ciscoasa(config-if)# duplex full</pre></p>	<p>Sets the duplex. 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.</p>
<p>Step 7</p> <pre>no shutdown</pre> <p>Example: <pre>ciscoasa(config-if)# no shutdown</pre></p>	<p>Enables the switch port. To disable the switch port, enter the shutdown command.</p>

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
<code>show interface</code>	Displays interface statistics.
<code>show interface ip brief</code>	Displays interface IP addresses and status.

Configuration Examples for ASA 5505 Interfaces

This section includes the following topics:

- [Access Port Example, page 10-11](#)
- [Trunk Port Example, page 10-12](#)

Access Port Example

The following example configures five VLAN interfaces, including the failover interface which is configured using the **failover lan** command:

```
ciscoasa(config)# interface vlan 100
ciscoasa(config-if)# nameif outside
ciscoasa(config-if)# security-level 0
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 200
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# ip address 10.2.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 300
ciscoasa(config-if)# nameif dmz
ciscoasa(config-if)# security-level 50
ciscoasa(config-if)# ip address 10.3.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 400
ciscoasa(config-if)# nameif backup-isp
ciscoasa(config-if)# security-level 50
ciscoasa(config-if)# ip address 10.1.2.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# failover lan faillink vlan500
ciscoasa(config)# failover interface ip faillink 10.4.1.1 255.255.255.0 standby 10.4.1.2
255.255.255.0

ciscoasa(config)# interface ethernet 0/0
ciscoasa(config-if)# switchport access vlan 100
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/1
ciscoasa(config-if)# switchport access vlan 200
```

```

ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/2
ciscoasa(config-if)# switchport access vlan 300
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/3
ciscoasa(config-if)# switchport access vlan 400
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/4
ciscoasa(config-if)# switchport access vlan 500
ciscoasa(config-if)# no shutdown

```

Trunk Port Example

The following example configures seven VLAN interfaces, including the failover interface which is configured using the **failover lan** command. VLANs 200, 201, and 202 are trunked on Ethernet 0/1.

```

ciscoasa(config)# interface vlan 100
ciscoasa(config-if)# nameif outside
ciscoasa(config-if)# security-level 0
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 200
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# ip address 10.2.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 201
ciscoasa(config-if)# nameif dept1
ciscoasa(config-if)# security-level 90
ciscoasa(config-if)# ip address 10.2.2.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 202
ciscoasa(config-if)# nameif dept2
ciscoasa(config-if)# security-level 90
ciscoasa(config-if)# ip address 10.2.3.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 300
ciscoasa(config-if)# nameif dmz
ciscoasa(config-if)# security-level 50
ciscoasa(config-if)# ip address 10.3.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 400
ciscoasa(config-if)# nameif backup-isp
ciscoasa(config-if)# security-level 50
ciscoasa(config-if)# ip address 10.1.2.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# failover lan faillink vlan500
ciscoasa(config)# failover interface ip faillink 10.4.1.1 255.255.255.0 standby 10.4.1.2
255.255.255.0

ciscoasa(config)# interface ethernet 0/0
ciscoasa(config-if)# switchport access vlan 100

```

```

ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/1
ciscoasa(config-if)# switchport mode trunk
ciscoasa(config-if)# switchport trunk allowed vlan 200-202
ciscoasa(config-if)# switchport trunk native vlan 5
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/2
ciscoasa(config-if)# switchport access vlan 300
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/3
ciscoasa(config-if)# switchport access vlan 400
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface ethernet 0/4
ciscoasa(config-if)# switchport access vlan 500
ciscoasa(config-if)# no shutdown

```

Where to Go Next

Complete the interface configuration according to [Chapter 11, “Completing Interface Configuration \(Routed Mode\)”](#) or [Chapter 12, “Completing Interface Configuration \(Transparent Mode\)”](#).

Feature History for ASA 5505 Interfaces

[Table 10-1](#) lists the release history for this feature.

Table 10-1 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 introduced the following command: switchport trunk native vlan .



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 11-1](#)
- [Licensing Requirements for Completing Interface Configuration in Routed Mode, page 11-2](#)
- [Guidelines and Limitations, page 11-5](#)
- [Default Settings, page 11-6](#)
- [Completing Interface Configuration in Routed Mode, page 11-6](#)
- [Turning Off and Turning On Interfaces, page 11-17](#)
- [Monitoring Interfaces, page 11-17](#)
- [Configuration Examples for Interfaces in Routed Mode, page 11-18](#)
- [Feature History for Interfaces in Routed Mode, page 11-19](#)



Note

For multiple context mode, complete the tasks in this section in the context execution space. Enter the **changeto context *name*** command to change to the context you want to configure.

Information About Completing Interface Configuration in Routed Mode

This section includes the following topics:

- [Security Levels, page 11-1](#)
- [Dual IP Stack \(IPv4 and IPv6\), page 11-2](#)

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 11-15](#) 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 11-15](#)), 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	<p>VLANs:</p> <p>Routed Mode:</p> <p>Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone)</p> <p>Security Plus License: 20</p> <p>Transparent Mode:</p> <p>Base License: 2 active VLANs in 1 bridge group.</p> <p>Security Plus License: 3 active VLANs: 2 active VLANs in 1 bridge group, and 1 active VLAN for the failover link.</p> <p>VLAN Trunks:</p> <p>Base License: None.</p> <p>Security Plus License: 8.</p>

Model	License Requirement
ASA 5510	<p>VLANs¹:</p> <p>Base License: 50</p> <p>Security Plus License: 100</p> <p>Interface Speed:</p> <p>Base License—All interfaces Fast Ethernet.</p> <p>Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet.</p> <p>Interfaces of all types²:</p> <p>Base License: 364</p> <p>Security Plus License: 564</p>
ASA 5520	<p>VLANs¹:</p> <p>Base License: 150.</p> <p>Interfaces of all types²:</p> <p>Base License: 764</p>
ASA 5540	<p>VLANs¹:</p> <p>Base License: 200</p> <p>Interfaces of all types²:</p> <p>Base License: 964</p>
ASA 5550	<p>VLANs¹:</p> <p>Base License: 400</p> <p>Interfaces of all types²:</p> <p>Base License: 1764</p>

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. For example:

```
interface gigabitethernet 0/0.100
vlan 100
```


2. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** command defined in the configuration counts against this limit. For example, both of the following interfaces count even if the GigabitEthernet 0/0 interface is defined as part of port-channel 1:

```
interface gigabitethernet 0/0
and
interface port-channel 1
```

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 9, “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 6-15](#).
- PPPoE is not supported in multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode. For transparent mode, see [Chapter 12, “Completing Interface Configuration \(Transparent Mode\)”](#).

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See [Chapter 7, “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 3-18](#).

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 11-10](#).

Completing Interface Configuration in Routed Mode

This section includes the following topics:

- [Task Flow for Completing Interface Configuration, page 11-7](#)
- [Configuring General Interface Parameters, page 11-7](#)
- [Configuring the MAC Address, MTU, and TCP MSS, page 11-10](#)
- [Configuring IPv6 Addressing, page 11-12](#)
- [Allowing Same Security Level Communication, page 11-15](#)

Task Flow for Completing Interface Configuration

-
- Step 1** Set up your interfaces depending on your model:
- ASA 5510 and higher—[Chapter 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15.](#)
- Step 3** (Multiple context mode) Enter the **changeto context name** command to change to the context you want to configure. Configure general interface parameters, including the interface name, security level, and IPv4 address. See the [“Configuring General Interface Parameters” section on page 11-7.](#)
- Step 4** (Optional) Configure the MAC address and the MTU. See the [“Configuring the MAC Address, MTU, and TCP MSS” section on page 11-10.](#)
- Step 5** (Optional) Configure IPv6 addressing. See the [“Configuring IPv6 Addressing” section on page 11-12.](#)
- 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 11-15.](#)
-

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 7, “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 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15.](#)
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505 or ASASM:</p> <pre>ciscoasa(config)# interface {vlan number mapped_name}</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabithethernet 0/0</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>nameif name</pre> <p>Example:</p> <pre>ciscoasa(config-if)# nameif inside</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value. Do not enter the no form, because that command causes all commands that refer to that name to be deleted.</p>
Step 3	Do one of the following:	

Command	Purpose
<p>ip address <i>ip_address</i> [<i>mask</i>] [standby <i>ip_address</i>]</p> <p>Example: ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0 standby 10.1.1.2</p>	<p>Sets the IP address manually.</p> <p>Note For use with failover, you must set the IP address and standby address manually; DHCP and PPPoE are not supported.</p> <p>The <i>ip_address</i> and <i>mask</i> arguments set the interface IP address and subnet mask.</p> <p>The standby <i>ip_address</i> argument is used for failover. See the “Configuring Active/Standby Failover” section on page 7-26 or the “Configuring Active/Active Failover” section on page 7-30 for more information.</p>
<p>ip address dhcp [setroute]</p> <p>Example: ciscoasa(config-if)# ip address dhcp</p>	<p>Obtains an IP address from a DHCP server.</p> <p>The setroute keyword lets the ASA use the default route supplied by the DHCP server.</p> <p>Reenter this command to reset the DHCP lease and request a new lease.</p> <p>If you do not enable the interface using the no shutdown command before you enter the ip address dhcp command, some DHCP requests might not be sent.</p>
<p>To obtain an IP address from a PPPoE server, see Chapter 9, “Configuring the PPPoE Client,” in the VPN configuration guide.</p>	<p>PPPoE is not supported in multiple context mode.</p>
<p>Step 4 security-level <i>number</i></p> <p>Example: ciscoasa(config-if)# security-level 50</p>	<p>Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest). See the “Security Levels” section on page 11-1.</p>
<p>Step 5 (Optional) management-only</p> <p>Example: ciscoasa(config-if)# management-only</p>	<p>Sets an interface to management-only mode so that it does not pass through traffic.</p> <p>By default, Management interfaces are configured as management-only. To disable this setting, enter the no management-only command.</p> <p>(ASA 5512-X through ASA 5555-X) You cannot disable management-only on the Management 0/0 interface.</p> <p>The management-only command is not supported for a redundant interface.</p>

Example

The following example configures parameters for VLAN 101:

```
ciscoasa(config)# interface vlan 101
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0
```

The following example configures parameters in multiple context mode for the context configuration. The interface ID is a mapped name.

```
ciscoasa/contextA(config)# interface int1
ciscoasa/contextA(config-if)# nameif outside
ciscoasa/contextA(config-if)# security-level 100
ciscoasa/contextA(config-if)# ip address 10.1.2.1 255.255.255.0
```

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 11-10.
- (Optional) Configure IPv6 addressing. See the [“Configuring IPv6 Addressing”](#) section on page 11-12.

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 6-3 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 6-24 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 9-8.

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—Chapter 9, “Starting Interface Configuration (ASA 5510 and Higher).”
 - ASA 5505—Chapter 10, “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 6-15.
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.
- To increase the MTU above 1500, enable jumbo frames on supported models according to the “Enabling Jumbo Frame Support (Supported Models)” section on page 9-35. Jumbo frames are supported by default on the ASASM; you do not need to enable them.

Detailed Steps

	Command	Purpose
Step 1	<p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505 or ASASM:</p> <pre>ciscoasa(config)# interface {vlan number mapped_name}</pre> <p>Example:</p> <pre>ciscoasa(config)# interface vlan 100</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>mac-address mac_address [standby mac_address]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# mac-address 000C.F142.4CDE</pre>	<p>Assigns a private MAC address to this interface. The <i>mac_address</i> is in H.H.H format, 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.</p> <p>The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.</p> <p>For use with failover, set the standby MAC address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.</p>

	Command	Purpose
Step 3	<pre>mtu interface_name bytes</pre> <p>Example: <pre>ciscoasa(config)# mtu inside 9200</pre></p>	<p>Sets the MTU between 300 and 65,535 bytes. The default is 1500 bytes.</p> <p>Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.</p> <p>For models that support jumbo frames, if you enter a value for any interface that is greater than 1500, then you need to enable jumbo frame support. See the “Enabling Jumbo Frame Support (Supported Models)” section on page 9-35.</p>
Step 4	<pre>sysopt connection tcpmss [minimum] bytes</pre> <p>Example: <pre>ciscoasa(config)# sysopt connection tcpmss 8500 ciscoasa(config)# sysopt connection tcpmss minimum 1290</pre></p>	<p>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 bytes to 0.</p> <p>For the minimum keyword, sets the maximum segment size to be no less than <i>bytes</i>, between 48 and 65535. The minimum feature is disabled by default (set to 0).</p>

What to Do Next

(Optional) Configure IPv6 addressing. See the [“Configuring IPv6 Addressing”](#) section on page 11-12.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the [“IPv6 Addresses”](#) section on page 49-5.

This section includes the following topics:

- [Information About IPv6, page 11-12](#)
- [Configuring a Global IPv6 Address, page 11-13](#)
- [Configuring IPv6 Neighbor Discovery, page 11-15](#)

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- [IPv6 Addressing, page 11-12](#)
- [Modified EUI-64 Interface IDs, page 11-13](#)

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.

**Note**

If you want to only configure the link-local addresses, see the **ipv6 enable** (to auto-configure) or **ipv6 address link-local** (to manually configure) command in the command reference.

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 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

Command	Purpose
<p>Step 1</p> <p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505 or ASASM:</p> <pre>ciscoasa(config)# interface {vlan number mapped_name}</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabithethernet 0/0</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
<p>Step 2</p> <p>Do one of the following:</p> <pre>ipv6 address autoconfig</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ipv6 address autoconfig</pre>	<p>Enables stateless autoconfiguration on the interface. Enabling stateless autoconfiguration on the interface configures IPv6 addresses based on 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.</p> <p>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 ipv6 nd suppress-ra command to suppress messages.</p>
<pre>ipv6 address ipv6-address/prefix-length [standby ipv6-address]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ipv6 address 2001:0DB8::BA98:0:3210/48</pre>	<p>Assigns a global address to the interface. When you assign a global address, the link-local address is automatically created for the interface.</p> <p>standby specifies the interface address used by the secondary unit or failover group in a failover pair.</p> <p>See the “IPv6 Addresses” section on page 49-5 for more information about IPv6 addressing.</p>

Command	Purpose
<p>ipv6 address <i>ipv6-prefix/prefix-length</i> <i>eui-64</i></p> <p>Example: ciscoasa(config-if)# ipv6 address 2001:0DB8::BA98::/48 eui-64</p>	<p>Assigns a global address to the interface by combining the specified prefix with an interface ID generated from the interface MAC address using the Modified EUI-64 format. When you assign a global address, the link-local address is automatically created for the interface.</p> <p>You do not need to specify the standby address; the interface ID will be generated automatically.</p> <p>See the “IPv6 Addresses” section on page 49-5 for more information about IPv6 addressing.</p>
<p>Step 3 (Optional)</p> <p>ipv6 enforce-eui64 <i>if_name</i></p> <p>Example: ciscoasa(config)# ipv6 enforce-eui64 inside</p>	<p>Enforces the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link.</p> <p>The <i>if_name</i> argument is the name of the interface, as specified by the nameif command, on which you are enabling the address format enforcement.</p> <p>See the “Modified EUI-64 Interface IDs” section on page 11-13 for more information.</p>

Configuring IPv6 Neighbor Discovery

See [Chapter 31, “Configuring IPv6 Neighbor Discovery,”](#) to configure IPv6 neighbor discovery.

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

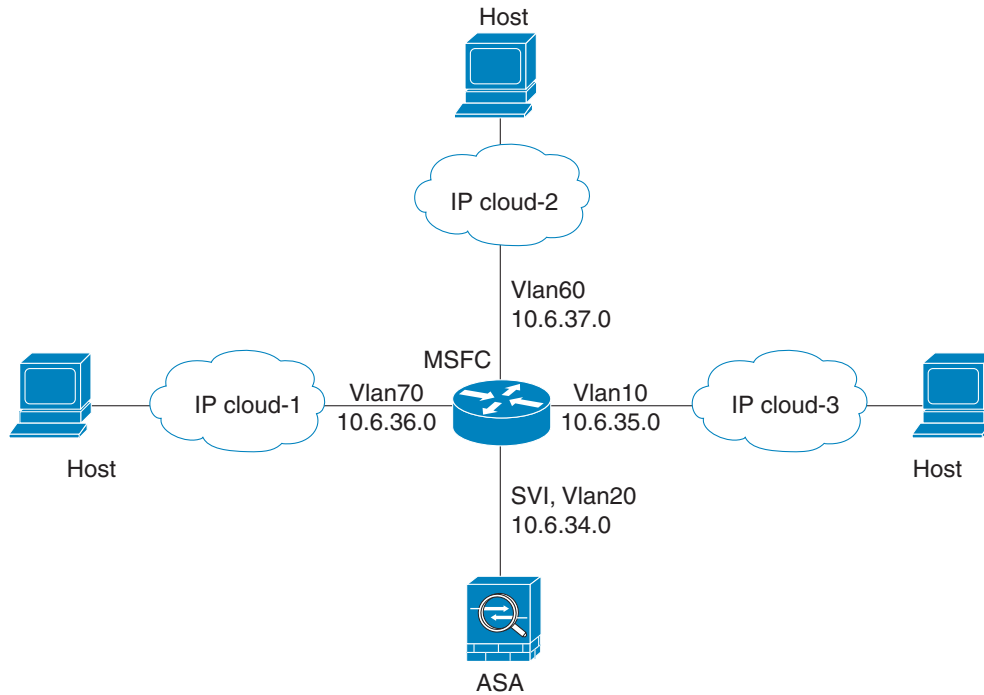
Intra-interface communication might be useful for VPN traffic that enters an interface, but is then routed out the same interface. The VPN traffic might be unencrypted in this case, or it might be reencrypted for another VPN connection. For example, if you have a hub and spoke VPN network, where the 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 11-1](#) shows a network where hosts on the same interface need to communicate.

Figure 11-1 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 11-1](#):

```
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

Command	Purpose
<code>same-security-traffic permit inter-interface</code>	Enables interfaces on the same security level so that they can communicate with each other.
<code>same-security-traffic permit intra-interface</code>	Enables communication between 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 reenables the interface within a context, only that context interface is affected. But if you disable or reenables the interface in the system execution space, then you affect that interface for all contexts.

Detailed Steps

	Command	Purpose
Step 1	<code>ciscoasa(config)# interface {vlan number mapped_name}</code> Example: <code>ciscoasa(config)# interface vlan 100</code>	If you are not already in interface configuration mode, enters interface configuration mode. In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.
Step 2	<code>shutdown</code> Example: <code>ciscoasa(config-if)# shutdown</code>	Disables the interface.
Step 3	<code>no shutdown</code> Example: <code>ciscoasa(config-if)# no shutdown</code>	Reenables the interface.

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
<code>show interface</code>	Displays interface statistics.
<code>show interface ip brief</code>	Displays interface IP addresses and status.

Configuration Examples for Interfaces in Routed Mode

This section includes the following topics:

- [ASA 5505 Example, page 11-18](#)

ASA 5505 Example

The following example configures three VLAN interfaces for the Base license. The third home interface cannot forward traffic to the business interface.

```
ciscoasa(config)# interface vlan 100
ciscoasa(config-if)# nameif outside
ciscoasa(config-if)# security-level 0
ciscoasa(config-if)# ip address dhcp
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 200
ciscoasa(config-if)# nameif business
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown

ciscoasa(config-if)# interface vlan 300
ciscoasa(config-if)# no forward interface vlan 200
ciscoasa(config-if)# nameif home
ciscoasa(config-if)# security-level 50
ciscoasa(config-if)# ip address 10.2.1.1 255.255.255.0
ciscoasa(config-if)# no shutdown
```

Feature History for Interfaces in Routed Mode

Table 11-1 lists the release history for this feature.

Table 11-1 Feature History for Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: <ul style="list-style-type: none"> • 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. Use the speed command to change the speed on the interface and use the show interface command to see what speed is currently configured for each interface.
Native VLAN support for the ASA 5505	7.2(4)/8.0(4)	You can now include the native VLAN in an ASA 5505 trunk port. We introduced the following command: switchport trunk native vlan .

Table 11-1 Feature History for Interfaces (continued)

Feature Name	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 introduced the following command: jumbo-frame reservation .
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 introduced the following command: flowcontrol .



Completing Interface Configuration (Transparent Mode)

This chapter includes tasks to complete the interface configuration for all models in transparent firewall mode.

This chapter includes the following sections:

- [Information About Completing Interface Configuration in Transparent Mode, page 12-1](#)
- [Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 12-3](#)
- [Guidelines and Limitations, page 12-5](#)
- [Default Settings, page 12-7](#)
- [Completing Interface Configuration in Transparent Mode, page 12-7](#)
- [Turning Off and Turning On Interfaces, page 12-19](#)
- [Monitoring Interfaces, page 12-19](#)
- [Configuration Examples for Interfaces in Transparent Mode, page 12-20](#)
- [Feature History for Interfaces in Transparent Mode, page 12-21](#)



Note

For multiple context mode, complete the tasks in this section in the context execution space. Enter the `changeto context name` command to change to the context you want to configure.

Information About Completing Interface Configuration in Transparent Mode

This section includes the following topics:

- [Bridge Groups in Transparent Mode, page 12-2](#)
- [Security Levels, page 12-2](#)

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 12-18 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 12-18), 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	<p>VLANs:</p> <p>Routed Mode:</p> <p>Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone)</p> <p>Security Plus License: 20</p> <p>Transparent Mode:</p> <p>Base License: 2 active VLANs in 1 bridge group.</p> <p>Security Plus License: 3 active VLANs: 2 active VLANs in 1 bridge group, and 1 active VLAN for the failover link.</p> <p>VLAN Trunks:</p> <p>Base License: None.</p> <p>Security Plus License: 8.</p>

Model	License Requirement
ASA 5510	<p>VLANs¹:</p> <p>Base License: 50</p> <p>Security Plus License: 100</p> <p>Interface Speed:</p> <p>Base License—All interfaces Fast Ethernet.</p> <p>Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet.</p> <p>Interfaces of all types²:</p> <p>Base License: 364</p> <p>Security Plus License: 564</p>
ASA 5520	<p>VLANs¹:</p> <p>Base License: 150.</p> <p>Interfaces of all types²:</p> <p>Base License: 764</p>
ASA 5540	<p>VLANs¹:</p> <p>Base License: 200</p> <p>Interfaces of all types²:</p> <p>Base License: 964</p>

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

- For an interface to count against the VLAN limit, you must assign a VLAN to it. For example:

```
interface gigabitethernet 0/0.100
  vlan 100
```
- The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** command defined in the configuration counts against this limit. For example, both of the following interfaces count even if the GigabitEthernet 0/0 interface is defined as part of port-channel 1:

```
interface gigabitethernet 0/0
and
interface port-channel 1
```

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 9, “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 using the **allocate-interface** command.

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 9-2](#) 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 12-8](#) 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 7, “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 3-18](#).

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 12-13](#).

Completing Interface Configuration in Transparent Mode

This section includes the following topics:

- [Task Flow for Completing Interface Configuration, page 12-8](#)
- [Configuring Bridge Groups, page 12-8](#)
- [Configuring General Interface Parameters, page 12-9](#)
- [Configuring a Management Interface \(ASA 5510 and Higher\), page 12-12](#)
- [Configuring the MAC Address, MTU, and TCP MSS, page 12-13](#)
- [Configuring IPv6 Addressing, page 12-16](#)
- [Allowing Same Security Level Communication, page 12-18](#)

Task Flow for Completing Interface Configuration

-
- Step 1** Set up your interfaces depending on your model:
- ASA 5510 and higher—[Chapter 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15.
- Step 3** (Multiple context mode) Enter the **changeto context** *name* command to change to the context you want to configure.
- Step 4** Configure one or more bridge groups, including the IPv4 address. See the [“Configuring Bridge Groups”](#) section on page 12-8.
- 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 12-9.
- 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 12-12.
- Step 7** (Optional) Configure the MAC address and the MTU. See the [“Configuring the MAC Address, MTU, and TCP MSS”](#) section on page 12-13.
- Step 8** (Optional) Configure IPv6 addressing. See the [“Configuring IPv6 Addressing”](#) section on page 12-16.
- Step 9** (Optional) Allow same security level communication, either by allowing communication between two interfaces or by allowing traffic to enter and exit the same interface. See the [“Allowing Same Security Level Communication”](#) section on page 12-18.
-

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

	Command	Purpose
Step 1	<pre>interface bvi <i>bridge_group_number</i></pre> <p>Example: <pre>ciscoasa(config)# interface bvi 1</pre></p>	Creates a bridge group, where <i>bridge_group_number</i> is an integer between 1 and 100.
Step 2	<pre>ip address <i>ip_address</i> [<i>mask</i>] [standby <i>ip_address</i>]</pre> <p>Example: <pre>ciscoasa(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2</pre></p>	<p>Specifies the management IP address for the bridge group.</p> <p>Do not assign a host address (/32 or 255.255.255.255) to the bridge group. 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. Therefore, 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.</p> <p>The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.</p> <p>The standby keyword and address is used for failover.</p>

Examples

The following example sets the management address and standby address of bridge group 1:

```
ciscoasa(config)# interface bvi 1  

ciscoasa(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2
```

What to Do Next

Configure general interface parameters. See the [“Configuring General Interface Parameters”](#) section on page 12-9.

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 12-12.

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 12-2](#).
- 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 7, “Configuring Failover,”](#) to configure the failover and state links.

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—[Chapter 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context *name*** command.

Detailed Steps

	Command	Purpose
Step 1	<p>For the ASA 5510 and higher:</p> <pre>interface {{redundant <i>number</i> port-channel <i>number</i> <i>physical_interface</i>}[.<i>subinterface</i>] <i>mapped_name</i>}</pre> <p>For the ASA 5505:</p> <pre>ciscoasa(config)# interface {vlan <i>number</i> <i>mapped_name</i>}</pre> <p>Example:</p> <pre>ciscoasa(config)# interface vlan 100</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant <i>number</i> argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel <i>number</i> argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID. Do not use this procedure for Management interfaces; see the “Configuring a Management Interface (ASA 5510 and Higher)” section on page 12-12 to configure the Management interface.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>bridge-group <i>number</i></pre> <p>Example:</p> <pre>ciscoasa(config-if)# bridge-group 1</pre>	<p>Assigns the interface to a bridge group, where <i>number</i> is an integer between 1 and 100. You can assign up to four interfaces to a bridge group. You cannot assign the same interface to more than one bridge group.</p>
Step 3	<pre>nameif <i>name</i></pre> <p>Example:</p> <pre>ciscoasa(config-if)# nameif inside</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value. Do not enter the no form, because that command causes all commands that refer to that name to be deleted.</p>
Step 4	<pre>security-level <i>number</i></pre> <p>Example:</p> <pre>ciscoasa(config-if)# security-level 50</pre>	<p>Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest).</p>

What to Do Next

- (Optional) Configure a management interface. See the “[Configuring a Management Interface \(ASA 5510 and Higher\)](#)” section on page 12-12.
- (Optional) Configure the MAC address and the MTU. See the “[Configuring the MAC Address, MTU, and TCP MSS](#)” section on page 12-13.
- (Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 12-16.

Configuring a Management Interface (ASA 5510 and Higher)

You can configure one management interface separate from the bridge group interfaces in single mode or per context. For more information, see the [“Management Interface” section on page 9-2](#).

Restrictions

- See the [“Management Interface” section on page 9-2](#).
- 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 9, “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 6-15](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<pre>interface {{port-channel number management slot/port}[.subinterface] mapped_name} Example: ciscoasa(config)# interface management 0/0.1</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode for the management interface.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1. The EtherChannel interface must have only Management member interfaces.</p> <p>Redundant interfaces do not support Management <i>slot/port</i> interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>nameif name Example: ciscoasa(config-if)# nameif management</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value. Do not enter the no form, because that command causes all commands that refer to that name to be deleted.</p>

	Command	Purpose
Step 3	Do one of the following:	
	<pre>ip address <i>ip_address</i> [<i>mask</i>] [standby <i>ip_address</i>]</pre> <p>Example: ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0 standby 10.1.1.2</p>	<p>Sets the IP address manually.</p> <p>Note For use with failover, you must set the IP address and standby address manually; DHCP is not supported.</p> <p>The <i>ip_address</i> and <i>mask</i> arguments set the interface IP address and subnet mask.</p> <p>The standby <i>ip_address</i> argument is used for failover. See the “Configuring Active/Standby Failover” section on page 7-26 or the “Configuring Active/Active Failover” section on page 7-30 for more information.</p>
	<pre>ip address dhcp [setroute]</pre> <p>Example: ciscoasa(config-if)# ip address dhcp</p>	<p>Obtains an IP address from a DHCP server.</p> <p>The setroute keyword lets the ASA use the default route supplied by the DHCP server.</p> <p>Reenter this command to reset the DHCP lease and request a new lease.</p> <p>If you do not enable the interface using the no shutdown command before you enter the ip address dhcp command, some DHCP requests might not be sent.</p>
Step 4	<pre>security-level <i>number</i></pre> <p>Example: ciscoasa(config-if)# security-level 50</p>	<p>Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest).</p>

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 12-13.
- (Optional) Configure IPv6 addressing. See the “Configuring IPv6 Addressing” section on page 12-16.

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 6-3](#) 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 6-24](#) 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 9-8](#).

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—[Chapter 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.
- To increase the MTU above 1500, enable jumbo frames on supported models according to the [“Enabling Jumbo Frame Support \(Supported Models\)” section on page 9-35](#). Jumbo frames are supported by default on the ASASM; you do not need to enable them.

Detailed Steps

Command	Purpose
<p>Step 1</p> <p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505 or ASASM:</p> <pre>ciscoasa(config)# interface {vlan number mapped_name}</pre> <p>Example:</p> <pre>ciscoasa(config)# interface vlan 100</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
<p>Step 2</p> <pre>mac-address mac_address [standby mac_address]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# mac-address 000C.F142.4CDE</pre>	<p>Assigns a private MAC address to this interface. The <i>mac_address</i> is in H.H.H format, 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.</p> <p>The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.</p> <p>For use with failover, set the standby MAC address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.</p>
<p>Step 3</p> <pre>mtu interface_name bytes</pre> <p>Example:</p> <pre>ciscoasa(config)# mtu inside 9200</pre>	<p>Sets the MTU between 300 and 65,535 bytes. The default is 1500 bytes.</p> <p>Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.</p> <p>For models that support jumbo frames, if you enter a value for any interface that is greater than 1500, then you need to enable jumbo frame support. See the “Enabling Jumbo Frame Support (Supported Models)” section on page 9-35.</p>
<p>Step 4</p> <pre>sysopt connection tcpmss [minimum] bytes</pre> <p>Example:</p> <pre>ciscoasa(config)# sysopt connection tcpmss 8500 ciscoasa(config)# sysopt connection tcpmss minimum 1290</pre>	<p>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 bytes to 0.</p> <p>For the minimum keyword, sets the maximum segment size to be no less than <i>bytes</i>, between 48 and 65535. The minimum feature is disabled by default (set to 0).</p>

What to Do Next

(Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 12-16.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the “[IPv6 Addresses](#)” section on page 49-5.

This section includes the following topics:

- [Information About IPv6, page 12-16](#)
- [Configuring a Global IPv6 Address, page 12-17](#)
- [Configuring IPv6 Neighbor Discovery, page 12-18](#)

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- [IPv6 Addressing, page 12-16](#)
- [Modified EUI-64 Interface IDs, page 12-16](#)
- [Unsupported Commands, page 12-17](#)

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 address 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.



Note

If you want to only configure the link-local addresses, see the **ipv6 enable** (to auto-configure) or **ipv6 address link-local** (to manually configure) command in the command reference.

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 9, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 10, “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 6-15](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<p>For the bridge group:</p> <pre>interface bvi bridge_group_id</pre> <p>For the management interface:</p> <pre>interface management_interface_id</pre> <p>Example: ciscoasa(config)# interface bvi 1</p>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p>
Step 2	<pre>ipv6 address ipv6-address/prefix-length [standby ipv6-address]</pre> <p>Example: ciscoasa(config-if)# ipv6 address 2001:0DB8::BA98:0:3210/48</p>	<p>Assigns a global address to the interface. When you assign a global address, the link-local address is automatically created for the interface (for a bridge group, for each member interface).</p> <p>standby specifies the interface address used by the secondary unit or failover group in a failover pair.</p> <p>Note The eui-64 keyword to use the Modified EUI-64 interface ID for the interface ID is not supported in transparent mode.</p> <p>See the “IPv6 Addresses” section on page 49-5 for more information about IPv6 addressing.</p>
Step 3	<p>(Optional)</p> <pre>ipv6 enforce-eui64 if_name</pre> <p>Example: ciscoasa(config)# ipv6 enforce-eui64 inside</p>	<p>Enforces the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link.</p> <p>The <i>if_name</i> argument is the name of the interface, as specified by the nameif command, on which you are enabling the address format enforcement.</p> <p>See the “Modified EUI-64 Interface IDs” section on page 12-16 for more information.</p>

Configuring IPv6 Neighbor Discovery

See Chapter 31, “Configuring IPv6 Neighbor Discovery,” to configure IPv6 neighbor discovery.

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

Command	Purpose
<code>same-security-traffic permit inter-interface</code>	Enables interfaces on the same security level so that they can communicate with each other.

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 reenables the interface within a context, only that context interface is affected. But if you disable or reenables the interface in the system execution space, then you affect that interface for all contexts.

Detailed Steps

	Command	Purpose
Step 1	<pre>ciscoasa(config)# interface {vlan number mapped_name}</pre> <p>Example: ciscoasa(config)# interface vlan 100 </p>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>shutdown</pre> <p>Example: ciscoasa(config-if)# shutdown </p>	Disables the interface.
Step 3	<pre>no shutdown</pre> <p>Example: ciscoasa(config-if)# no shutdown </p>	Reenables the interface.

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
<code>show interface</code>	Displays interface statistics.
<code>show interface ip brief</code>	Displays interface IP addresses and status.
<code>show bridge-group</code>	Shows bridge group information.

Configuration Examples for Interfaces in Transparent Mode

The following example includes two bridge groups of three interfaces each, plus a management-only interface:

```
interface gigabitethernet 0/0
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 0/1
  nameif outside1
  security-level 0
  bridge-group 1
  no shutdown
interface gigabitethernet 0/2
  nameif dmz1
  security-level 50
  bridge-group 1
  no shutdown
interface bvi 1
  ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2

interface gigabitethernet 1/0
  nameif inside2
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/1
  nameif outside2
  security-level 0
  bridge-group 2
  no shutdown
interface gigabitethernet 1/2
  nameif dmz2
  security-level 50
  bridge-group 2
  no shutdown
interface bvi 2
  ip address 10.3.5.8 255.255.255.0 standby 10.3.5.9

interface management 0/0
  nameif mgmt
  security-level 100
  ip address 10.2.1.1 255.255.255.0 standby 10.2.1.2
  no shutdown
```

Feature History for Interfaces in Transparent Mode

Table 12-1 lists each feature change and the platform release in which it was implemented.

Table 12-1 Feature History for Interfaces in Transparent Mode

Feature Name	Platform Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: <ul style="list-style-type: none"> • 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. Use the speed command to change the speed on the interface and use the show interface command to see what speed is currently configured for each interface.
Native VLAN support for the ASA 5505	7.2(4)/8.0(4)	You can now include the native VLAN in an ASA 5505 trunk port. We introduced the following command: switchport trunk native vlan .

Table 12-1 Feature History for Interfaces in Transparent Mode (continued)

Feature Name	Platform Releases	Feature Information
Jumbo packet support for the ASA 5580	8.1(1)	<p>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.</p> <p>We introduced the following command: jumbo-frame reservation.</p>
Increased VLANs for the ASA 5580	8.1(2)	<p>The number of VLANs supported on the ASA 5580 are increased from 100 to 250.</p>
IPv6 support for transparent mode	8.2(1)	<p>IPv6 support was introduced for transparent firewall mode.</p>
Support for Pause Frames for Flow Control on the ASA 5580 10-Gigabit Ethernet Interfaces	8.2(2)	<p>You can now enable pause (XOFF) frames for flow control.</p> <p>We introduced the following command: flowcontrol.</p>
Bridge groups for transparent mode	8.4(1)	<p>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.</p> <p>We introduced the following commands: interface bvi, show bridge-group.</p>



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 13-1](#)
- [Setting the Date and Time, page 13-4](#)
- [Configuring the Master Passphrase, page 13-8](#)
- [Configuring the DNS Server, page 13-13](#)
- http://www.cisco.com/en/US/products/ps6121/products_tech_note09186a0080aaeff5.shtml, page 13-14
- [Performing Password Recovery, page 13-14](#)
- [Monitoring DNS Cache, page 13-16](#)

Configuring the Hostname, Domain Name, and Passwords

This section includes the following topics:

- [Setting the Login Password, page 13-2](#)
- [Changing the Enable Password, page 13-3](#)
- [Setting the Hostname, page 13-3](#)
- [Setting the Domain Name, page 13-4](#)
- [Feature History for the Hostname, Domain Name, and Passwords, page 13-4](#)

Setting the Login Password

The login password is used for Telnet access when you do not configure Telnet authentication (see the [“Configuring Authentication for CLI and ASDM Access”](#) section on page 41-20). You also use this password when accessing the ASASM from the switch with the **session** command.

Prerequisites

Enable Telnet access according to the [“Configuring Telnet Access”](#) section on page 41-3.

To set the login password, enter the following command:

Command	Purpose
{ passwd password } <i>password</i> [encrypted]	<p>Sets the login password. 9.1(1): The default password is “cisco.” 9.1(2) and later: There is no default password.</p> <p>You can enter passwd or password. The password is a case-sensitive password of up to 16 alphanumeric and special characters. You can use any character in the password except a question mark or a space.</p> <p>The password is saved in the configuration in encrypted form, so you cannot view the original password after you enter it. If for some reason you need to copy the password to another ASA but do not know the original password, you can enter the passwd command with the encrypted password and the encrypted keyword. Normally, you only see this keyword when you enter the show running-config passwd command.</p> <p>Use the no password command to restore the password to the default setting; in 9.1(2) and later, the no form of the command removes the password.</p>

Changing the Enable Password

The enable password lets you enter privileged EXEC mode if you do not configure enable authentication (see the “[Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\)](#)” section on page 41-21).

The enable password also lets you log into ASDM with a blank username if you do not configure HTTP authentication (see the “[Configuring Authentication for CLI and ASDM Access](#)” section on page 41-20).

To change the enable password, enter the following command:

Command	Purpose
<p><code>enable password password</code></p> <p>Example: <pre>hostname(config)# passwd Pa\$\$w0rd</pre></p>	<p>Changes the enable password. By default, the enable password is blank.</p> <p>The <i>password</i> argument is a case-sensitive password of up to 16 alphanumeric and special characters. You can use any character in the password except a question mark or a space.</p> <p>This command changes the password for the highest privilege level. If you configure local command authorization, you can set enable passwords for each privilege level from 0 to 15.</p> <p>The password is saved in the configuration in encrypted form, so you cannot view the original password after you enter it. Enter the enable password command without a password to set the password to the default, which is blank.</p>

Setting the Hostname

When you set a hostname for the ASA, that name appears in the command line prompt. If you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands.

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, but can be used by the **banner** command **\$(hostname)** token.

To set the hostname, enter the following command:

Command	Purpose
<p><code>hostname name</code></p> <p>Example: <pre>asa(config)# hostname farscape farscape(config)#</pre></p>	<p>Specifies the hostname for the ASA or for a context. The default hostname is “asa.”</p> <p>This name can be up to 63 characters. The hostname must start and end with a letter or digit, and have only letters, digits, or a hyphen.</p>

Setting the Domain Name

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.”

For multiple context mode, you can set the domain name for each context, as well as within the system execution space.

To set the domain name, enter the following command:

Command	Purpose
<code>domain-name name</code>	Specifies the domain name for the ASA.
Example: <code>ciscoasa(config)# domain-name example.com</code>	The default domain name is default.domain.invalid.

Feature History for the Hostname, Domain Name, and Passwords

Table 13-1 lists each feature change and the platform release in which it was implemented.

Table 13-1 Feature History for the Master Passphrase

Feature Name	Platform Releases	Feature Information
Removal of the default Telnet password	9.0(2), 9.1(2)	<p>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 (the aaa authentication telnet console command).</p> <p>Formerly, when you cleared the password, the ASA restored the default of “cisco.” Now when you clear the password, the password is removed.</p> <p>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.</p> <p>We modified the following command: passwd.</p>

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 Time Zone and Daylight Saving Time Date Range, page 13-5](#)
- [Setting the Date and Time Using an NTP Server, page 13-6](#)
- [Setting the Date and Time Manually, page 13-8](#)

Setting the Time Zone and Daylight Saving Time Date Range

To set the time zone and daylight saving time date range, perform the following steps:

	Command	Purpose
Step 1	<pre>clock timezone zone [-]hours [minutes]</pre> <p>Example: <pre>ciscoasa(config)# clock timezone PST -8</pre></p>	<p>Sets the time zone. By default, the time zone is UTC and the daylight saving time date range is from 2:00 a.m. on the first Sunday in April to 2:00 a.m. on the last Sunday in October.</p> <p>Where <i>zone</i> specifies the time zone as a string, for example, PST for Pacific Standard Time.</p> <p>The <i>[-]hours</i> value sets the number of hours of offset from UTC. For example, PST is -8 hours.</p> <p>The <i>minutes</i> value sets the number of minutes of offset from UTC.</p>
Step 2	To change the date range for daylight saving time from the default, enter one of the following commands. The default recurring date range is from 2:00 a.m. on the second Sunday in March to 2:00 a.m. on the first Sunday in November.	

Command	Purpose
<pre>clock summer-time zone date {day month month day} year hh:mm {day month month day} year hh:mm [offset]</pre> <p>Example:</p> <pre>ciscoasa(config)# clock summer-time PDT 1 April 2010 2:00 60</pre>	<p>Sets the start and end dates for daylight saving time as a specific date in a specific year. If you use this command, you need to reset the dates every year.</p> <p>The <i>zone</i> value specifies the time zone as a string, for example, PDT for Pacific Daylight Time.</p> <p>The <i>day</i> value sets the day of the month, from 1 to 31. You can enter the day and month as April 1 or as 1 April, for example, depending on your standard date format.</p> <p>The <i>month</i> value sets the month as a string. You can enter the day and month as April 1 or as 1 April, depending on your standard date format.</p> <p>The <i>year</i> value sets the year using four digits, for example, 2004. The year range is 1993 to 2035.</p> <p>The <i>hh:mm</i> value sets the hour and minutes in 24-hour time.</p> <p>The <i>offset</i> value sets the number of minutes to change the time for daylight saving time. By default, the value is 60 minutes.</p>
<pre>clock summer-time zone recurring [week weekday month hh:mm week weekday month hh:mm] [offset]</pre> <p>Example:</p> <pre>ciscoasa(config)# clock summer-time PDT recurring first Monday April 2:00 60</pre>	<p>Specifies the start and end dates for daylight saving time, in the form of a day and time of the month, and not a specific date in a year.</p> <p>This command enables you to set a recurring date range that you do not need to change yearly.</p> <p>The <i>zone</i> value specifies the time zone as a string, for example, PDT for Pacific Daylight Time.</p> <p>The <i>week</i> value specifies the week of the month as an integer between 1 and 4 or as the words first or last. For example, if the day might fall in the partial fifth week, then specify last.</p> <p>The <i>weekday</i> value specifies the day of the week: Monday, Tuesday, Wednesday, and so on.</p> <p>The <i>month</i> value sets the month as a string.</p> <p>The <i>hh:mm</i> value sets the hour and minutes in 24-hour time.</p> <p>The <i>offset</i> value sets the number of minutes to change the time for daylight savings time. By default, the value is 60 minutes.</p>

Setting the Date and Time Using an NTP Server

To obtain the date and time from an NTP server, perform the following steps :

Detailed Steps

	Command	Purpose
Step 1	<pre>ntp authenticate</pre> <p>Example:</p> <pre>ciscoasa(config)# ntp authenticate</pre>	Enables authentication with an NTP server.

<p>Step 2</p>	<pre>ntp trusted-key key_id</pre> <p>Example: ciscoasa(config)# ntp trusted-key 1 </p>	<p>Specifies an authentication key ID to be a trusted key, which is required for authentication with an NTP server.</p> <p>The <i>key_id</i> argument is a value between 1 and 4294967295. You can enter multiple trusted keys for use with multiple servers.</p>
<p>Step 3</p>	<pre>ntp authentication-key key_id md5 key</pre> <p>Example: hostname(config)# ntp authentication-key 1 md5 aNiceKey </p>	<p>Sets a key to authenticate with an NTP server.</p> <p>The <i>key_id</i> argument is the ID you set in Step 2 using the ntp trusted-key command, and the <i>key</i> argument is a string up to 32 characters long.</p>
<p>Step 4</p>	<pre>ntp server ip_address [key key_id] [source interface_name] [prefer]</pre> <p>Example: hostname(config)# ntp server 10.1.1.1 key 1 prefer </p>	<p>Identifies an NTP server.</p> <p>The <i>key_id</i> argument is the ID you set in Step 2 using the ntp trusted-key command.</p> <p>The source interface_name keyword-argument pair identifies the outgoing interface for NTP packets if you do not want to use the default interface in the routing table. Because the system does not include any interfaces in multiple context mode, specify an interface name defined in the admin context.</p> <p>The prefer keyword sets this NTP server as the preferred server if multiple servers have similar accuracy. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If servers are of similar accuracy, then the prefer keyword specifies which of those servers to use. However, if a server is significantly more accurate than the preferred one, the ASA uses the more accurate one. For example, the ASA uses a server of stratum 2 over a server of stratum 3 that is preferred.</p> <p>You can identify multiple servers; the ASA uses the most accurate server.</p> <p>Note In multiple context mode, set the time in the system configuration only.</p>

Setting the Date and Time Manually

To set the date and time manually, enter the following command:

Command	Purpose
<pre>clock set hh:mm:ss {month day day month} year</pre> <p>Example: hostname# clock set 20:54:00 april 1 2004</p>	<p>Sets the date time manually.</p> <p>The <i>hh:mm:ss</i> argument sets the hour, minutes, and seconds in 24-hour time. For example, enter 20:54:00 for 8:54 pm.</p> <p>The <i>day</i> value sets the day of the month, from 1 to 31. You can enter the day and month as april 1 or as 1 april, for example, depending on your standard date format.</p> <p>The <i>month</i> value sets the month. Depending on your standard date format, you can enter the day and month as april 1 or as 1 april.</p> <p>The <i>year</i> value sets the year using four digits, for example, 2004. The year range is from 1993 to 2035.</p> <p>The default time zone is UTC. If you change the time zone after you enter the clock set command using the clock timezone command, the time automatically adjusts to the new time zone.</p> <p>This command sets the time in the hardware chip, and does not save the time in the configuration file. This time endures reboots. Unlike the other clock commands, this command is a privileged EXEC command. To reset the clock, you need to set a new time with the clock set command.</p>

Configuring the Master Passphrase

This section includes the following topics:

- [Information About the Master Passphrase, page 13-8](#)
- [Licensing Requirements for the Master Passphrase, page 13-9](#)
- [Guidelines and Limitations, page 13-9](#)
- [Adding or Changing the Master Passphrase, page 13-9](#)
- [Disabling the Master Passphrase, page 13-11](#)
- [Recovering the Master Passphrase, page 13-12](#)
- [Feature History for the Master Passphrase, page 13-13](#)

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.

Adding or Changing the Master Passphrase

This procedure will only be accepted in a secure session, for example by console, SSH, or ASDM via HTTPS.

To add or change the master passphrase, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<p>key config-key password-encryption [<i>new_passphrase</i> [<i>old_passphrase</i>]]</p> <p>Example: ciscoasa(config)# key config-key password-encryption Old key: bumblebee New key: haverford Confirm key: haverford</p>	<p>Sets the passphrase used for generating the encryption key. The passphrase must be between 8 and 128 characters long. All characters except a backspace and double quotes are accepted for the passphrase.</p> <p>If you do not enter the new passphrase in the command, you are prompted for it.</p> <p>To change the passphrase, you must enter the old passphrase.</p> <p>See the “Examples” section on page 13-11 for examples of the interactive prompts.</p> <p>Note Use the interactive prompts to enter passwords to avoid having the passwords logged in the command history buffer.</p> <p>Use the no key config-key password-encrypt command with caution, because it changes the encrypted passwords into plain text passwords. You can use the no form of this command when downgrading to a software version that does not support password encryption.</p>
Step 2	<p>password encryption aes</p> <p>Example: ciscoasa(config)# password encryption aes</p>	<p>Enables password encryption. As soon as password encryption is enabled and the master passphrase is available, all the user passwords will be encrypted. The running configuration will show the passwords in the encrypted format.</p> <p>If the passphrase is not configured at the time that password encryption is enabled, the command will succeed in anticipation that the passphrase will be available in the future.</p> <p>If you later disable password encryption using the no password encryption aes command, 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.</p>
Step 3	<p>write memory</p> <p>Example: ciscoasa(config)# write memory</p>	<p>Saves the runtime value of the master passphrase and the resulting configuration. If you do not enter this command, passwords in startup configuration may still be visible if they were not saved with encryption previously.</p> <p>In addition, in multiple context mode the master passphrase is changed in the system context configuration. As a result, the passwords in all contexts will be affected. If the write memory command is not entered in the system context mode, but not in all user contexts, then the encrypted passwords in user contexts may be stale. Alternatively, use the write memory all command in the system context to save all configurations.</p>

Examples

The following example shows that no previous key was present:

```
hostname (config)# key config-key password-encryption 12345678
```

The following example shows that a key already exists:

```
Hostname (config)# key config-key password-encryption 23456789  
Old key: 12345678  
hostname (config)#
```

In the following example, you want to key in interactively, but a key already exists. The Old key, New key, and Confirm key prompts appear on your screen if you enter the **key config-key password-encryption** command and press **Enter** to access interactive mode.

```
hostname (config)# key config-key password-encryption  
Old key: 12345678  
New key: 23456789  
Confirm key: 23456789
```

In the following example, you want to key in interactively, but no key is present. The New key and Confirm key prompts appear on your screen if you are in interactive mode.

```
hostname (config)# key config-key password-encryption  
New key: 12345678  
Confirm key: 12345678
```

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 13-12.

This procedure works only in a secure session; that is, by Telnet, SSH, or ASDM via HTTPS.

To disable the master passphrase, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<p>no key config-key password-encryption [old_passphrase]</p> <p>Example: ciscoasa(config)# no key config-key password-encryption</p> <p>Warning! You have chosen to revert the encrypted passwords to plain text. This operation will expose passwords in the configuration and therefore exercise caution while viewing, storing, and copying configuration.</p> <p>Old key: bumblebee</p>	<p>Removes the master passphrase.</p> <p>If you do not enter the passphrase in the command, you are prompted for it.</p>
Step 2	<p>write memory</p> <p>Example: ciscoasa(config)# write memory</p>	<p>Saves the runtime value of the master passphrase and the resulting configuration. The non-volatile memory containing the passphrase will be erased and overwritten with the 0xFF pattern.</p> <p>In multiple mode, the master passphrase is changed in the system context configuration. As a result, the passwords in all contexts will be affected. If the write memory command is not entered in the system context mode, but not in all user contexts, then the encrypted passwords in user contexts may be stale. Alternatively, use the write memory all command in the system context to save all configurations.</p>

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:

	Command	Purpose
Step 1	<p>write erase</p> <p>Example: ciscoasa(config)# write erase</p>	Removes the master key and the configuration that includes the encrypted passwords.
Step 2	<p>reload</p> <p>Example: ciscoasa(config)# reload</p>	Reloads the ASA with the startup configuration, without any master key or encrypted passwords.

Feature History for the Master Passphrase

Table 13-2 lists each feature change and the platform release in which it was implemented.

Table 13-2 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 commands: key config-key password-encryption , password encryption aes , clear configure password encryption aes , show running-config password encryption aes , show password encryption .
Password Encryption Visibility	8.4(1)	We modified the show password encryption command.

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 example, most commands require you to enter an IP address and can only use a name when you manually configure the **name** command to associate a name with an IP address and enable use of the names using the **names** command.

For information about dynamic DNS, see the “[Configuring DDNS](#)” section on page 15-2.

Prerequisites

Make sure that you configure the appropriate routing for any interface on which you enable DNS domain lookup so you can reach the DNS server. See the “[Information About Routing](#)” section on page 24-1 for more information about routing.

To configure the DNS server, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<code>dns domain-lookup interface_name</code> Example: hostname(config)# dns domain-lookup inside	Enables the ASA to send DNS requests to a DNS server to perform a name lookup for supported commands.
Step 2	<code>dns server-group DefaultDNS</code> Example: hostname(config)# dns server-group DefaultDNS	Specifies the DNS server group that the ASA uses for outgoing requests. Other DNS server groups can be configured for VPN tunnel groups. See the tunnel-group command in the command reference for more information.
Step 3	<code>name-server ip_address [ip_address2] [...] [ip_address6]</code> Example: hostname(config-dns-server-group)# name-server 10.1.1.5 192.168.1.67 209.165.201.6	Specifies one or more DNS servers. You can enter all six IP addresses in the same command, separated by spaces, or you can enter each command separately. The ASA tries each DNS server in order until it receives a response.

http://www.cisco.com/en/US/products/ps6121/products_tech_note09186a0080aaeff5.shtml

Performing Password Recovery

This section includes the following topics:

- [Recovering Passwords for the ASA, page 13-14](#)
- [Disabling Password Recovery, page 13-16](#)

Recovering Passwords for the ASA

To recover passwords for the ASA, perform the following steps:

-
- Step 1** Connect to the ASA console port according to the instructions in “[Accessing the ASA Services Module Command-Line Interface](#)” section on page 3-2 or the “[Accessing the Appliance Command-Line Interface](#)” section on page 3-1.
 - Step 2** Power off the ASA, and then power it on.
 - Step 3** After startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 To update the configuration register value, enter the following command:

```
rommon #1> confreg 0x41
Update Config Register (0x41) in NVRAM...
```

Step 5 To set the ASA to ignore the startup configuration, enter the following command:

```
rommon #1> confreg
```

The ASA displays the current configuration register value, and asks whether you want to change it:

```
Current Configuration Register: 0x00000041
Configuration Summary:
  boot default image from Flash
  ignore system configuration

Do you wish to change this configuration? y/n [n]: y
```

Step 6 Record the current configuration register value, so you can restore it later.

Step 7 At the prompt, enter **Y** to change the value.

The ASA prompts you for new values.

Step 8 Accept the default values for all settings, except for the "disable system configuration?" value.

Step 9 At the prompt, enter **Y**.

Step 10 Reload the ASA by entering the following command:

```
rommon #2> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading disk0:/asa800-226-k8.bin... Booting...Loading...
```

The ASA loads the default configuration instead of the startup configuration.

Step 11 Access the privileged EXEC mode by entering the following command:

```
ciscoasa# enable
```

Step 12 When prompted for the password, press **Enter**.

The password is blank.

Step 13 Load the startup configuration by entering the following command:

```
ciscoasa# copy startup-config running-config
```

Step 14 Access the global configuration mode by entering the following command:

```
ciscoasa# configure terminal
```

Step 15 Change the passwords, as required, in the default configuration by entering the following commands:

```
ciscoasa(config)# password password
ciscoasa(config)# enable password password
ciscoasa(config)# username name password password
```

Step 16 Load the default configuration by entering the following command:

```
ciscoasa(config)# no config-register
```

The default configuration register value is 0x1. For more information about the configuration register, see the command reference.

Step 17 Save the new passwords to the startup configuration by entering the following command:

```
ciscoasa(config)# copy running-config startup-config
```

Disabling Password Recovery

To disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the ASA, enter the following command:

Command	Purpose
<code>no service password-recovery</code>	Disables password recovery.
Example: <pre>ciscoasa (config)# no service password-recovery</pre>	

On the ASA, the **no service password-recovery** command prevents you from entering ROMMON mode with the configuration intact. When you enter ROMMON mode, the ASA prompts you to erase all Flash file systems. You cannot enter ROMMON mode without first performing this erasure. If you choose not to erase the Flash file system, the ASA reloads. Because password recovery depends on using ROMMON mode and maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to restore the system to an operating state, load a new image and a backup configuration file, if available.

The **service password-recovery** command appears in the configuration file for information only. When you enter the command at the CLI prompt, the setting is saved in NVRAM. The only way to change the setting is to enter the command at the CLI prompt. Loading a new configuration with a different version of the command does not change the setting. If you disable password recovery when the ASA is configured to ignore the startup configuration at startup (in preparation for password recovery), then the ASA changes the setting to load the startup configuration as usual. If you use failover, and the standby unit is configured to ignore the startup configuration, then the same change is made to the configuration register when the **no service password recovery** command replicates to the standby unit.

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, enter the following command:

Command	Purpose
<code>show dns-hosts</code>	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 14-1](#)
- [Licensing Requirements for DHCP, page 14-2](#)
- [Guidelines and Limitations, page 14-2](#)
- [Configuring DHCP Services, page 14-4](#)
- [Additional References, page 14-11](#)
- [Monitoring DHCP Services, page 14-11](#)
- [Feature History for DHCP Services, page 14-12](#)

Information About DHCP Services

- [Information About the DHCP Server, page 14-1](#)
- [Information About the DHCP Relay Agent, page 14-2](#)

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 14-1 shows the licensing requirements for DHCP.

Table 14-1 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 14-4 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 DHCP 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 14-4](#)
- [Configuring the DHCP Relay Agent, page 14-8](#)

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 14-5](#)
- [Configuring DHCP Options, page 14-6](#)

Enabling the DHCP Server

To enable the DHCP server on an ASA interface, perform the following steps.

Detailed Steps

	Command	Purpose
Step 1	<p>dhcpcd address <i>ip_address if_name</i></p> <p>Example: ciscoasa(config)# dhcpcd address 10.0.1.101-10.0.1.110 inside</p>	<p>Creates a DHCP address pool. The ASA assigns a client one of the addresses from this pool to use for a given period of time. These addresses are the local, untranslated addresses for the directly connected network.</p> <p>The address pool must be on the same subnet as the ASA interface.</p>
Step 2	<p>dhcpcd dns <i>dns1 [dns2]</i></p> <p>Example: ciscoasa(config)# dhcpcd dns 209.165.201.2 209.165.202.129</p>	(Optional) Specifies the IP address(es) of the DNS server(s).
Step 3	<p>dhcpcd wins <i>wins1 [wins2]</i></p> <p>Example: ciscoasa(config)# dhcpcd wins 209.165.201.5</p>	(Optional) Specifies the IP address(es) of the WINS server(s). You can specify up to two WINS servers.
Step 4	<p>dhcpcd lease <i>lease_length</i></p> <p>Example: ciscoasa(config)# dhcpcd lease 3000</p>	(Optional) Changes the lease length to be granted to the client. The lease length equals the amount of time in seconds that the client can use its allocated IP address before the lease expires. Enter a value from 0 to 1,048,575. The default value is 3600 seconds.
Step 5	<p>dhcpcd domain <i>domain_name</i></p> <p>Example: ciscoasa(config)# dhcpcd domain example.com</p>	(Optional) Configures the domain name.
Step 6	<p>dhcpcd ping_timeout <i>milliseconds</i></p> <p>Example: ciscoasa(config)# dhcpcd ping timeout 20</p>	(Optional) Configures the DHCP ping timeout value for ICMP packets. To avoid address conflicts, the ASA sends two ICMP ping packets to an address before assigning that address to a DHCP client.
Step 7	<p>dhcpcd option 3 ip <i>gateway_ip</i></p> <p>Example: ciscoasa(config)# dhcpcd option 3 ip 10.10.1.1</p>	Defines a default gateway that is sent to DHCP clients. If you do not use the dhcpcd option 3 command to define the default gateway, DHCP clients use the ASA interface IP address that is closest to the DHCP clients by default; the ASA does not use the management interface IP address. As a result, the DHCP ACK does not include this option.
Step 8	<p>dhcpcd enable <i>interface_name</i></p> <p>Example: ciscoasa(config)# dhcpcd enable outside</p>	Enables the DHCP daemon within the ASA to listen for DHCP client requests on the enabled interface.

Configuring DHCP Options

The ASA supports the DHCP options listed in RFC 2132 to send information. This section includes the following topics:

- [Options that Return an IP Address, page 14-6](#)
- [Options that Return a Text String, page 14-6](#)
- [Options that Return a Hexadecimal Value, page 14-6](#)

Options that Return an IP Address

Command	Purpose
<code>dhcpd option code ip addr_1 [addr_2]</code>	Configures a DHCP option that returns one or two IP addresses.
Example: <pre>ciscoasa(config)# dhcpd option 2 ip 10.10.1.1 10.10.1.2</pre>	

Options that Return a Text String

Command	Purpose
<code>dhcpd option code ascii text</code>	Configures a DHCP option that returns a text string.
Example: <pre>ciscoasa(config)# dhcpd option 2 ascii examplestring</pre>	

Options that Return a Hexadecimal Value

Command	Purpose
<code>dhcpd option code hex value</code>	Configures a DHCP option that returns a hexadecimal value.
Example: <pre>ciscoasa(config)# dhcpd option 2 hex 22.0011.01.FF1111.00FF.0000.AAAA.1111.1111 .1111.11</pre>	



Note

The ASA does not verify that the option type and value that you provide match the expected type and value for the option code as defined in RFC 2132. For example, you can enter the **dhcpd option 46 ascii hello** command, and the ASA accepts the configuration, although option 46 is defined in RFC 2132 to expect a single-digit, hexadecimal value. For more information about option codes and their associated types and expected values, see RFC 2132.

Table 14-2 shows the DHCP options that are not supported by the **dhcpd option** command.

Table 14-2 **Unsupported DHCP Options**

Option Code	Description
0	DHCPOPT_PAD
1	HCPOPT_SUBNET_MASK
12	DHCPOPT_HOST_NAME
50	DHCPOPT_REQUESTED_ADDRESS
51	DHCPOPT_LEASE_TIME
52	DHCPOPT_OPTION_OVERLOAD
53	DHCPOPT_MESSAGE_TYPE
54	DHCPOPT_SERVER_IDENTIFIER
58	DHCPOPT_RENEWAL_TIME
59	DHCPOPT_REBINDING_TIME
61	DHCPOPT_CLIENT_IDENTIFIER
67	DHCPOPT_BOOT_FILE_NAME
82	DHCPOPT_RELAY_INFORMATION
255	DHCPOPT_END

DHCP options 3, 66, and 150 are used to configure Cisco IP phones. For more information about configuring these options, see the [“Using Cisco IP Phones with a DHCP Server”](#) section on page 14-7.

Using Cisco IP Phones with a DHCP Server

Cisco IP phones download their configuration from a TFTP server. When a Cisco IP phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information.

- DHCP option 150 provides the IP addresses of a list of TFTP servers.
- DHCP option 66 gives the IP address or the hostname of a single TFTP server.



Note

Cisco IP phones can also include DHCP option 3 in their requests, which sets the default route.

A single request might include both options 150 and 66. In this case, the ASA DHCP server provides values for both options in the response if they are already configured on the ASA.

To send information to use for any option number, enter the following command:

Command	Purpose
<code>dhcpcd option number value</code>	Provides information for DHCP requests that include an option number as specified in RFC 2132.
Example: <code>ciscoasa(config)# dhcpcd option 2</code>	

To send information to use for option 66, enter the following command:

Command	Purpose
<code>dhcpd option 66 ascii server_name</code>	Provides the IP address or name of a TFTP server for option 66.
Example: <pre>ciscoasa(config)# dhcpd option 66 ascii exampleserver</pre>	

To send information to use for option 150, enter the following command:

Command	Purpose
<code>dhcpd option 150 ip server_ip1 [server_ip2]</code>	Provides the IP address or names of one or two TFTP servers for option 150. The <i>server_ip1</i> is the IP address or name of the primary TFTP server while <i>server_ip2</i> is the IP address or name of the secondary TFTP server. A maximum of two TFTP servers can be identified using option 150.
Example: <pre>ciscoasa(config)# dhcpd option 150 ip 10.10.1.1</pre>	

To send information to use for option 3, enter the following command:

Command	Purpose
<code>dhcpd option 3 ip router_ip1</code>	Sets the default route.
Example: <pre>ciscoasa(config)# dhcpd option 3 ip 10.10.1.1</pre>	

Configuring the DHCP Relay Agent

- [Configuring the DHCPv4 Relay Agent, page 14-8](#)
- [Configuring the DHCPv6 Relay Agent, page 14-10](#)

Configuring the DHCPv4 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

	Command	Purpose
Step 1	Do one or both of the following: For a global server: dhcprelay server <i>ip_address if_name</i> Example: ciscoasa(config)# dhcprelay server 209.165.201.5 outside ciscoasa(config)# dhcprelay server 209.165.201.8 outside ciscoasa(config)# dhcprelay server 209.165.202.150 it	Specifies a global DHCP server IP address and the interface through which it is reachable.
	For an interface-specific server: interface <i>interface_id</i> dhcprelay server <i>ip_address</i> Example: ciscoasa(config)# interface gigabitethernet 0/0 ciscoasa(config)# dhcprelay server 209.165.201.6 ciscoasa(config)# dhcprelay server 209.165.201.7 ciscoasa(config)# interface gigabitethernet 0/1 ciscoasa(config)# dhcprelay server 209.165.202.155 ciscoasa(config)# dhcprelay server 209.165.202.156	Specifies the interface ID connected to the DHCP client network, and the DHCP server IP address to be used for DHCP requests that enter that interface. Note that you do not specify the egress interface for the requests, as in the global dhcprelay server command; instead, the ASA uses the routing table to determine the egress interface.
Step 2	dhcprelay enable <i>interface</i> Example: ciscoasa(config)# dhcprelay enable inside ciscoasa(config)# dhcprelay enable dmz ciscoasa(config)# dhcprelay enable eng1 ciscoasa(config)# dhcprelay enable eng2 ciscoasa(config)# dhcprelay enable mktg	Enables the DHCP relay service on the interface connected to the DHCP clients. You can enable DHCP relay on multiple interfaces.
Step 3	dhcprelay timeout <i>seconds</i> Example: ciscoasa(config)# dhcprelay timeout 25	(Optional) Sets the number of seconds allowed for DHCP relay address handling.
Step 4	dhcprelay setroute <i>interface_name</i> Example: ciscoasa(config)# dhcprelay setroute inside	(Optional) Changes the first default router address in the packet sent from the DHCP server to the address of the ASA interface. 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.
Step 5	(Optional) Do one of the following:	

Command	Purpose
<pre>interface interface_id dhcprelay information trusted</pre> <p>Example:</p> <pre>ciscoasa(config)# interface gigabitethernet 0/0 ciscoasa(config-if)# dhcprelay information trusted</pre>	<p>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.</p>
<pre>dhcprelay information trust-all</pre> <p>Example:</p> <pre>ciscoasa(config)# dhcprelay information trust-all</pre>	<p>Configures all client interfaces as trusted.</p>

Configuring the DHCPv6 Relay Agent

When a DHCPv6 request enters an interface, then the ASA relays the request to all DHCPv6 global servers.

Detailed Steps

	Command	Purpose
Step 1	<pre>ipv6 dhcprelay server ipv6_address [interface]</pre> <p>Example:</p> <pre>ciscoasa(config)# ipv6 dhcprelay server 3FFB:C00:C18:6:A8BB:CCFF:FE03:2701</pre>	<p>Specifies the IPv6 DHCP server destination address to which client messages are forwarded.</p> <p>The <i>ipv6-address</i> argument can be a link-scoped unicast, multicast, site-scoped unicast, or global IPv6 address. Unspecified, loopback, and node-local multicast addresses are not allowed as the relay destination. The optional <i>interface</i> argument specifies the egress interface for a destination. Client messages are forwarded to the destination address through the link to which the egress interface is connected. If the specified address is a link-scoped address, then you must specify the interface.</p>
Step 2	<pre>ipv6 dhcprelay enable interface</pre> <p>Example:</p> <pre>ciscoasa(config)# ipv6 dhcprelay enable inside</pre>	<p>Enables DHCPv6 relay service on a client interface.</p>
Step 3	<pre>ipv6 dhcprelay timeout seconds</pre> <p>Example:</p> <pre>ciscoasa(config)# ipv6 dhcprelay timeout 25</pre>	<p>(Optional) Specifies the amount of time in seconds that is allowed for responses from the DHCPv6 server to pass to the DHCPv6 client through the relay binding for relay address handling.</p> <p>Valid values for the <i>seconds</i> argument range from 1 to 3600. The default is 60 seconds.</p>

Additional References

For additional information related to implementing DHCPv6, see the following section:

- [RFCs, page 14-11](#)

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, enter one or more of the following commands:

Command	Purpose
<code>show running-config dhcpd</code>	Shows the current DHCP configuration.
<code>show running-config dhcprelay</code>	Shows the current DHCP relay service status.
Tools > Command Line Interface Enter the <code>show ip address dhcp lease proxy</code> command, then click Send	Shows proxy entries in the IPL table.
Tools > Command Line Interface Enter the <code>show ip address dhcp lease summary</code> command, then click Send .	Shows summary for the entry.
Tools > Command Line Interface Enter the <code>show ip address dhcp lease server</code> command, then click Send .	Shows server entries in the IPL table.
<code>show ipv6 dhcprelay binding</code>	Shows the relay binding entries that were created by the relay agent.
<code>show ipv6 dhcprelay statistics</code>	Shows DHCP relay agent statistics for IPv6.
<code>clear config ipv6 dhcprelay</code>	Clears the IPv6 DHCP relay configuration.

Feature History for DHCP Services

Table 14-3 each feature change and the platform release in which it was implemented.

Table 14-3 Feature History for DHCP Services

Feature Name	Releases	Description
DHCP	7.0(1)	<p>The ASA can provide a DHCP server or DHCP relay services to DHCP clients attached to ASA interfaces.</p> <p>We introduced the following commands: dhcp client update dns, dhcpd address, dhcpd domain, dhcpd enable, dhcpd lease, dhcpd option, dhcpd ping timeout, dhcpd update dns, dhcpd wins, dhcp-network-scope, dhcprelay enable, dhcprelay server, dhcprelay setroute, dhcp-server, show running-config dhcpd, and show running-config dhcprelay.</p>
DHCP for IPv6 (DHCPv6)	9.0(1)	<p>Support for IPv6 was added.</p> <p>We introduced the following commands: ipv6 dhcprelay server, ipv6 dhcprelay enable, ipv6 dhcprelay timeout, clear config ipv6 dhcprelay, ipv6 nd managed-config-flag, ipv6 nd other-config-flag, debug ipv6 dhcp, debug ipv6 dhcprelay, show ipv6 dhcprelay binding, clear ipv6 dhcprelay binding, show ipv6 dhcprelay statistics, and clear ipv6 dhcprelay statistics.</p>
DHCP relay servers per interface (IPv4 only)	9.1(2)	<p>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.</p> <p>We introduced or modified the following commands: dhcprelay server (interface config mode), clear configure dhcprelay, show running-config dhcprelay.</p>
DHCP lease information	9.1(4)	<p>You can now view DHCP Client Lease information</p> <p>We introduced the following screen: Monitoring > Interfaces > DHCP> DHCP Lease Information.</p>
DHCP trusted interfaces	9.1(2)	<p>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.</p> <p>We introduced or modified the following commands: dhcprelay information trusted, dhcprelay information trust-all, show running-config dhcprelay.</p>

Table 14-3 Feature History for DHCP Services

Feature Name	Releases	Description
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 15-1](#)
- [Licensing Requirements for DDNS, page 15-2](#)
- [Guidelines and Limitations, page 15-2](#)
- [Configuring DDNS, page 15-2](#)
- [Configuration Examples for DDNS, page 15-3](#)
- [DDNS Monitoring Commands, page 15-9](#)
- [Feature History for DDNS, page 15-9](#)

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 13-13](#). To configure DHCP, see the [“Configuring the DHCP Server” section on page 14-4](#).

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 DDNS

This section describes examples for configuring the ASA to support Dynamic DNS. When you use DHCP and dynamic DNS update, this configures a host automatically for network access whenever it attaches to the IP network. You can locate and reach the host using its permanent, unique DNS hostname. Mobile hosts, for example, can move freely without user or administrator intervention.

DDNS 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 is held 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 two most common DDNS update configurations are the following:

- The DHCP client updates the A RR, while the DHCP server updates the PTR RR.
- The DHCP server updates both the A RR and PTR RR.

In general, the DHCP server maintains DNS PTR RRs on behalf of clients. Clients may be configured to perform all desired DNS updates. The server may be configured to honor these updates or not. To update the PTR RR, the DHCP server must know the FQDN of the client. The client provides an FQDN to the server using a DHCP option called Client FQDN.

Configuration Examples for DDNS

The following examples present five common scenarios:

- [Example 1: Client Updates Both A and PTR RRs for Static IP Addresses, page 15-3](#)
- [Example 2: Client Updates Both A and PTR RRs; DHCP Server Honors Client Update Request; FQDN Provided Through Configuration, page 15-4](#)
- [Example 3: Client Includes FQDN Option Instructing Server Not to Update Either RR; Server Overrides Client and Updates Both RRs., page 15-5](#)
- [Example 4: Client Asks Server To Perform Both Updates; Server Configured to Update PTR RR Only; Honors Client Request and Updates Both A and PTR RR, page 15-6](#)
- [Example 5: Client Updates A RR; Server Updates PTR RR, page 15-7](#)

Example 1: Client Updates Both A and PTR RRs for Static IP Addresses

The following example shows how to configure the client to request that it update both A and PTR resource records for static IP addresses.

To configure this scenario, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	ddns update method <i>name</i> 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 <i>mapped_name</i> Example: ciscoasa(DDNS-update-method)# interface eth1	Configures an interface eth1 and enters interface configuration mode.

	Command	Purpose
Step 4	<pre>ddns update [method-name hostname hostname]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ddns update ddns-2 ciscoasa(config-if)# ddns update hostname asa.example.com</pre>	Associates the the DDNS method ddns-2 with the eth1 interface and an update hostname.
Step 5	<pre>ip address ip_address [mask] [standby ip_address]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ip address 10.0.0.40 255.255.255.0</pre>	Configures a static IP address for the interace eth1.

Example 2: Client Updates Both A and PTR RRs; DHCP Server Honors Client Update Request; FQDN Provided Through Configuration

The following example shows how to configure the DHCP client to request that it update both the A and PTR RRs, and the DHCP server to honor these requests.

To configure this scenario, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>dhcp-client update dns [server {both none}]</pre> <p>Example:</p> <pre>ciscoasa(config)# dhcp-client update dns server none</pre>	Configures the DHCP client to request that the DHCP server perform no updates
Step 2	<pre>ddns update method name</pre> <p>Example:</p> <pre>ciscoasa(config)# ddns update method ddns-2</pre>	Creates a DDNS update method ddns-2 that dynamically updates DNS resource records (RRs)
Step 3	<pre>ddns both</pre> <p>Example:</p> <pre>ciscoasa(DDNS-update-method)# ddns both</pre>	Specifies that the client updates both the DNS A and PTR resource records (RRs).
Step 4	<pre>interface mapped_name</pre> <p>Example:</p> <pre>ciscoasa(DDNS-update-method)# interface Ethernet0</pre>	Configures an interface Ethernet 0 and enters interface configuration mode.

	Command	Purpose
Step 5	ddns update [<i>method-name</i> hostname <i>hostname</i>] Example: <pre>ciscoasa(config-if)# ddns update ddns-2 ciscoasa(config-if)# ddns update hostname asa.example.com</pre>	Associates the DDNS method ddns-2 with the Ethernet0 interface and an update hostname.
Step 6	ip address dhcp Example: <pre>ciscoasa(if-config)# ip address dhcp</pre>	Uses DHCP to obtain an IP address for the interface.
Step 7	dhcpd update dns [both] [override] [interface <i>srv_ifc_name</i>] Example: <pre>ciscoasa(if-config)# dhcpd update dns</pre>	Configures DHCP server to perform DDNS updates.

Example 3: Client Includes FQDN Option Instructing Server Not to Update Either RR; Server Overrides Client and Updates Both RRs.

The following example shows how to configure the DHCP client to include the FQDN option that instruct the DHCP server not to honor either the A or PTR updates. The example also shows how to configure the server to override the client request. As a result, the client does not perform any updates.

To configure this scenario, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	ddns update method <i>name</i> Example: <pre>ciscoasa(config)# ddns update method ddns-2</pre>	Creates a DDNS update method ddns-2 that dynamically updates DNS resource records (RRs).
Step 2	ddns both Example: <pre>ciscoasa(DDNS-update-method)# ddns both</pre>	Specifies that the client updates both the DNS A and PTR resource records (RRs).
Step 3	interface <i>mapped_name</i> Example: <pre>ciscoasa(DDNS-update-method)# interface Ethernet0</pre>	Configures an interface Ethernet 0 and enters interface configuration mode.

	Command	Purpose
Step 4	<pre>ddns update [method-name hostname hostname]</pre> <p>Example:</p> <pre>ciscoasa(config-if)# ddns update ddns-2 ciscoasa(config-if)# ddns update hostname asa.example.com</pre>	Associates the the DDNS method ddns-2 with the Ethernet0 interface and an update hostname.
Step 5	<pre>dhcp-client update dns [server {both none}]</pre> <p>Example:</p> <pre>ciscoasa(config)# dhcp-client update dns server none</pre>	Configures the DHCP client to request that the DHCP server perform no updates.
Step 6	<pre>ip address dhcp</pre> <p>Example:</p> <pre>ciscoasa(if-config)# ip address dhcp</pre>	Uses DHCP to obtain an IP address for the interface.
Step 7	<pre>dhcpcd update dns [both] [override] [interface srv_ifc_name]</pre> <p>Example:</p> <pre>ciscoasa(if-config)# dhcpcd update dns both override</pre>	Configures DHCP server to override the client update requests.

Example 4: Client Asks Server To Perform Both Updates; Server Configured to Update PTR RR Only; Honors Client Request and Updates Both A and PTR RR

The following example shows how to configure the server to perform only PTR RR updates by default. However, the server honors the client request that it perform both A and PTR updates. The server also forms the FQDN by appending the domain name (example.com) to the hostname that the client (asa) has provided.

To configure this scenario, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	interface <i>mapped_name</i> 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 [<i>method-name</i> hostname <i>hostname</i>] Example: ciscoasa(config-if)# ddns update hostname asa	Configures the DHCP client on interface Ethernet 0.
Step 4	dhcpd update dns [both] [override] [interface <i>srv_ifc_name</i>] Example: ciscoasa(config-if)# dhcpd update dns	Configures DHCP server to perform DDNS updates.
Step 5	dhcpd domain <i>domain_name</i> [interface <i>if_name</i>] Example: ciscoasa(config-if)# dhcpd domain example.com	Defines the DNS domain name for DHCP clients.

Example 5: Client Updates A RR; Server Updates PTR RR

The following example shows how to configure the client to update the A resource record and how to configure the server to update the PTR records. Also, the client uses the domain name from the DHCP server to form the FQDN.

To configure this scenario, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	ddns update method <i>name</i> 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 <i>mapped_name</i> Example: ciscoasa(DDNS-update-method)# interface Ethernet0	Configures an interface Ethernet 0.
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 [<i>method-name</i> hostname <i>hostname</i>] 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 <i>srv_ifc_name</i>] Example: ciscoasa(if-config)# dhcpd update dns	Configures DHCP server to perform DDNS updates.
Step 7	dhcpd domain <i>domain_name</i> [interface <i>if_name</i>] Example: ciscoasa(config-if)# dhcpd domain example.com	Defines the DNS domain name for DHCP clients.

DDNS Monitoring Commands

To monitor DDNS, enter one of the following commands:

Command	Purpose
<code>show running-config ddns</code>	Shows the current DDNS configuration.
<code>show running-config dns server-group</code>	Shows the current DNS server group status.

Feature History for DDNS

[Table 15-1](#) lists each feature change and the platform release in which it was implemented.

Table 15-1 Feature History for DDNS

Feature Name	Releases	Feature Information
DDNS	7.0(1)	We introduced this feature. We introduced the following commands: ddns , ddns update , dhcp client update dns , dhcpd update dns , show running-config ddns , and show running-config dns server-group .



CHAPTER 16

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 16-1](#)
- [Guidelines and Limitations, page 16-1](#)
- [Licensing Requirements for WCCP, page 16-2](#)
- [Enabling WCCP Redirection, page 16-3](#)
- [WCCP Monitoring Commands, page 16-4](#)
- [Feature History for WCCP, page 16-4](#)

Information About WCCP

Web Cache Communication Protocol (WCCP) is a content routing protocol 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 through 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.

Enabling WCCP Redirection



Note

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 redirection is supported only on the ingress of an interface. The only topology that the ASA supports is when client and cache engine are behind the same interface of the ASA and the cache engine can directly communicate with the client, without going through the ASA.

The following configuration tasks assume you have already installed and configured the cache engines that you want to include in your network.

To configure WCCP redirection, perform the following steps:

	Command	Purpose
Step 1	<pre>wccp {web-cache service_number} [redirect-list access_list] [group-list access_list] [password password]</pre> <p>Example: hostname (config)# wccp web-cache</p>	<p>Enables a WCCP service group and identifies the service to be redirected. (Optional) Also defines which cache engines can participate in the service group, and what traffic should be redirected to the cache engine.</p> <p>The standard service is web-cache, which intercepts TCP port 80 (HTTP) traffic and redirects that traffic to the cache engines, but you can identify a service number (if desired) between 0 and 254. For example, to transparently redirect native FTP traffic to a cache engine, use WCCP service 60. You can enter this command multiple times for each service group that you want to enable.</p> <p>The redirect-list <i>access_list</i> argument controls traffic that is redirected to this service group.</p> <p>The group-list <i>access_list</i> argument determines which web cache IP addresses are allowed to participate in the service group.</p> <p>The password <i>password</i> argument specifies MD5 authentication for messages that are received from the service group. Messages that are not accepted by the authentication are discarded.</p>
Step 2	<pre>wccp interface interface_name {web-cache service_number} redirect in</pre> <p>Example: hostname (config)# wccp interface inside web-cache redirect in</p>	<p>Identifies an interface and enables WCCP redirection on the interface.</p> <p>The standard service is web-cache, which intercepts TCP port 80 (HTTP) traffic and redirects that traffic to the cache engines, but you can identify a service number (if desired) between 0 and 254. For example, to transparently redirect native FTP traffic to a cache engine, use WCCP service 60. You can enter this command multiple times for each service group that you want to enable.</p>

Examples

For example, to enable the standard web-cache service and redirect HTTP traffic that enters the inside interface to a web cache, enter the following commands:

```
hostname (config)# wccp web-cache
hostname (config)# wccp interface inside web-cache redirect in
```

WCCP Monitoring Commands

To monitor WCCP, enter one of the following commands:

Command	Purpose
<code>show running-config wccp</code>	Shows the current WCCP configuration.
<code>show running-config wccp interface</code>	Shows the current WCCP interfaces status.

Feature History for WCCP

[Table 16-1](#) lists the release history for this feature.

Table 16-1 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 commands: wccp and wccp interface



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 17-1](#)
- [Licensing Requirements for Objects, page 17-1](#)
- [Configuring Objects, page 17-2](#)
- [Monitoring Objects, page 17-19](#)
- [Feature History for Objects, page 17-19](#)

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 17-2](#)
- [Configuring Service Objects and Service Groups, page 17-5](#)
- [Configuring Local User Groups, page 17-11](#)
- [Configuring Security Group Object Groups, page 17-13](#)
- [Configuring Regular Expressions, page 17-14](#)
- [Configuring Time Ranges, page 17-18](#)

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 17-2](#)
- [Configuring a Network Object Group, page 17-3](#)

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

	Command	Purpose
Step 1	<p>object network <i>obj_name</i></p> <p>Example: ciscoasa(config)# object-network OBJECT1</p>	<p>Creates a new network object. The <i>obj_name</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters:</p> <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” <p>The prompt changes to network object configuration mode.</p>
Step 2	<p>{host <i>ip_addr</i> subnet <i>net_addr net_mask</i> range <i>ip_addr_1 ip_addr_2</i> fqdn <i>fully_qualified_domain_name</i>}</p> <p>Example: ciscoasa(config-network-object)# host 10.2.2.2</p>	<p>Assigns the IP address or FQDN to the named object.</p> <p>Note You cannot configure NAT for an FQDN object.</p>
Step 3	<p>description <i>text</i></p> <p>Example: ciscoasa(config-network-object)# description Engineering Network</p>	<p>Adds a description to the object.</p>

Examples

To create a network object, enter the following commands:

```
hostname (config)# object network OBJECT1
hostname (config-network-object)# host 10.2.2.2
```

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

	Command	Purpose
Step 1	object-group network <i>grp_id</i> Example: ciscoasa(config)# object-group network admins	Adds a network group. The <i>grp_id</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters: <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” The prompt changes to protocol configuration mode.
Step 2	description <i>text</i> Example: ciscoasa(config-network)# Administrator Addresses	(Optional) Adds a description. The description can be up to 200 characters.
Step 3	Add one or more of the following group members:	
	network-object object <i>name</i> Example: ciscoasa(config-network)# network-object host 10.2.2.4	Adds an object to the network object group.
	network-object { host <i>ipv4_address</i> <i>ipv4_address mask</i> <i>ipv6-address/prefix-length</i> } Example: ciscoasa(config-network)# network-object host 10.2.2.4	Adds a host or network inline, either IPv4 or IPv6.
	group-object <i>group_id</i> Example: ciscoasa(config-network)# group-object Engineering_groups	Adds an existing object group under this object group. The nested group must be of the same type.

Example

To create a network group that includes the IP addresses of three administrators, enter the following commands:

```
hostname (config)# object-group network admins
hostname (config-protocol)# description Administrator Addresses
hostname (config-protocol)# network-object host 10.2.2.4
hostname (config-protocol)# network-object host 10.2.2.78
hostname (config-protocol)# network-object host 10.2.2.34
```

Create network object groups for privileged users from various departments by entering the following commands:

```
hostname (config)# object-group network eng
hostname (config-network)# network-object host 10.1.1.5
```

```
hostname (config-network)# network-object host 10.1.1.9
hostname (config-network)# network-object host 10.1.1.89

hostname (config)# object-group network hr
hostname (config-network)# network-object host 10.1.2.8
hostname (config-network)# network-object host 10.1.2.12

hostname (config)# object-group network finance
hostname (config-network)# network-object host 10.1.4.89
hostname (config-network)# network-object host 10.1.4.100
```

You then nest all three groups together as follows:

```
hostname (config)# object-group network admin
hostname (config-network)# group-object eng
hostname (config-network)# group-object hr
hostname (config-network)# group-object finance
```

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 17-5](#)
- [Configuring a Service Group, page 17-6](#)
- [Configuring a TCP or UDP Port Service Group, page 17-8](#)
- [Configuring an ICMP Group, page 17-10](#)
- [Configuring an ICMP Group, page 17-10](#)

Configuring a Service Object

The service object can contain a protocol, ICMP, ICMPv6, TCP or UDP port or port ranges.

Detailed Steps

	Command	Purpose
Step 1	<p>object service <i>obj_name</i></p> <p>Example: <pre>ciscoasa(config)# object-service SERVOBJECT1</pre></p>	<p>Creates a new service object. The <i>obj_name</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters:</p> <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” <p>The prompt changes to service object configuration mode.</p>
Step 2	<p>service {<i>protocol</i> {icmp icmp6} [<i>icmp-type</i> [<i>icmp_code</i>]] {tcp udp} [source <i>operator port</i>] [destination <i>operator port</i>]}</p> <p>Example: <pre>ciscoasa(config-service-object)# service tcp destination eq www</pre></p>	<p>Adds a service definition to the object.</p> <p>The <i>protocol</i> argument specifies an IP protocol name or number. If you specify the icmp, icmp6, tcp, or udp protocols, you can specify additional properties; for other protocols, there are no additional properties.</p> <p>For icmp and icmp6 (for ICMP version 6), you can optionally include the ICMP type, such as “echo” or the type number. If you specify a type, you can optionally include an ICMP code, between 1 and 255.</p> <p>For TCP and UDP, you can optionally specify the source and/or destination ports, between 0 and 65535. For a list of supported names, see the CLI help. The operator can be:</p> <ul style="list-style-type: none"> • eq—Equals the port number. • gt—Greater than the port number. • lt—Less than the port number. • neq—Not equal to the port number. • range—A range of ports. Specify two numbers separated by a space, such as range 1024 4500.

Example

To create a service object, enter the following commands:

```
hostname (config)# object service SERVOBJECT1
hostname (config-service-object)# service tcp destination eq ssh
```

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

Command	Purpose
<p>Step 1</p> <p>object-group service <i>grp_id</i></p> <p>Example: <pre>ciscoasa(config)# object-group service services1</pre></p>	<p>Adds a service group. The <i>grp_id</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters:</p> <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” <p>The prompt changes to service configuration mode.</p>
<p>Step 2</p> <p>Add one or more of the following group members:</p>	
<p>service-object <i>protocol</i></p> <p>Example: <pre>ciscoasa(config-service-object-group)# service-object ipsec</pre></p>	<p>Identifies the protocol name or number, between 0 and 255.</p>
<p>service-object {tcp udp tcp-udp} [source <i>operator number</i>] [destination <i>operator number</i>]</p> <p>Example: <pre>ciscoasa(config-service-object-group)# service-object tcp destination eq domain</pre></p>	<p>You can specify the source and/or destination ports, between 0 and 65535, for the TCP and UDP protocols; tcp-udp matches both protocols. For a list of supported names, see the CLI help. The operator can be:</p> <ul style="list-style-type: none"> • eq—Equals the port number. • gt—Greater than the port number. • lt—Less than the port number. • neq—Not equal to the port number. • range—A range of ports. Specify two numbers separated by a space, such as range 1024 4500.
<p>service-object {icmp icmp6} [<i>icmp_type</i>] [<i>icmp_code</i>]</p> <p>Example: <pre>ciscoasa(config-service-object-group)# service-object icmp echo</pre></p>	<p>Specifies that the service type is for ICMP or ICMPv6 connections. You can optionally specify the ICMP type by name or number, between 0 and 255.</p> <p>The optional <i>icmp_code</i> specifies an ICMP code, between 1 and 255.</p>
<p>service-object object <i>name</i></p> <p>Example: <pre>ciscoasa(config-service-object-group)# service-object object my-service</pre></p>	<p>Specifies a service object name, created with the object service command.</p>

Command	Purpose
group-object <i>group_id</i> Example: ciscoasa(config-service-object-group)# group-object Engineering_groups	Adds an existing object group under this object group. The nested group must be of the same type.
Step 3 description <i>text</i> Example: ciscoasa(config-service-object-group)# description DNS Group	(Optional) Adds a description. The description can be up to 200 characters.

Examples

The following example shows how to add both TCP and UDP services to a service object group:

```
hostname(config)# object-group service CommonApps
hostname(config-service-object-group)# service-object tcp destination eq ftp
hostname(config-service-object-group)# service-object tcp-udp destination eq www
hostname(config-service-object-group)# service-object tcp destination eq h323
hostname(config-service-object-group)# service-object tcp destination eq https
hostname(config-service-object-group)# service-object udp destination eq ntp
```

The following example shows how to add multiple service objects to a service object group:

```
hostname(config)# service object SSH
hostname(config-service-object)# service tcp destination eq ssh
hostname(config)# service object EIGRP
hostname(config-service-object)# service eigrp
hostname(config)# service object HTTPS
hostname(config-service-object)# service tcp source range 0 1024 destination eq https
hostname(config)# object-group service Group1
hostname(config-service-object-group)# service-object object SSH
hostname(config-service-object-group)# service-object object EIGRP
hostname(config-service-object-group)# service-object object HTTPS
```

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).

	Command	Purpose
Step 1	<p>object-group service <i>grp_id</i> {tcp udp tcp-udp}</p> <p>Example: ciscoasa(config)# object-group service services1 tcp-udp</p>	<p>Adds a service group.</p> <p>The object keyword adds an additional object to the service object group.</p> <p>The <i>grp_id</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters:</p> <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” <p>Specifies the protocol for the services (ports) you want to add with either the tcp, udp, or tcp-udp keywords. Enter the tcp-udp keyword if your service uses both TCP and UDP with the same port number, for example, DNS (port53).</p> <p>The prompt changes to service configuration mode.</p>
Step 2	<p>Add one or more of the following group members:</p> <p>port-object {eq <i>port</i> range <i>begin_port</i> <i>end_port</i>}</p> <p>Example: ciscoasa(config-service)# port-object eq domain</p> <p>group-object <i>group_id</i></p> <p>Example: ciscoasa(config-network)# group-object Engineering_groups</p>	<p>Defines the ports in the group. Enter the command for each port or range of ports. For a list of permitted keywords and well-known port assignments, see the “Protocols and Applications” section on page 49-11.</p> <p>Adds an existing object group under this object group. The nested group must be of the same type.</p>
Step 3	<p>description <i>text</i></p> <p>Example: ciscoasa(config-service)# description DNS Group</p>	<p>(Optional) Adds a description. The description can be up to 200 characters.</p>

Example

To create service groups that include DNS (TCP/UDP), LDAP (TCP), and RADIUS (UDP), enter the following commands:

```
hostname (config)# object-group service services1 tcp-udp
hostname (config-service)# description DNS Group
hostname (config-service)# port-object eq domain

hostname (config)# object-group service services2 udp
hostname (config-service)# description RADIUS Group
hostname (config-service)# port-object eq radius
hostname (config-service)# port-object eq radius-acct

hostname (config)# object-group service services3 tcp
hostname (config-service)# description LDAP Group
hostname (config-service)# port-object eq ldap
```

Configuring an ICMP Group

An ICMP group includes multiple ICMP types.

Detailed Steps

	Command	Purpose
Step 1	object-group icmp-type <i>grp_id</i> Example: ciscoasa(config)# object-group icmp-type ping	Adds an ICMP type object group. The <i>grp_id</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters: <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” The prompt changes to ICMP type configuration mode.
Step 2	Add one or more of the following group members: icmp-object <i>icmp-type</i> Example: ciscoasa(config-icmp-type)# icmp-object echo-reply group-object <i>group_id</i> Example: ciscoasa(config-network)# group-object Engineering_groups	Defines the ICMP types in the group. Enter the command for each type. For a list of ICMP types, see the “ICMP Types” section on page 49-15 . Adds an existing object group under this object group. The nested group must be of the same type.
Step 3	description <i>text</i> Example: ciscoasa(config-icmp-type)# description Ping Group	(Optional) Adds a description. The description can be up to 200 characters.

Example

Create an ICMP type group that includes echo-reply and echo (for controlling ping) by entering the following commands:

```
hostname (config)# object-group icmp-type ping
hostname (config-service)# description Ping Group
hostname (config-service)# icmp-object echo
hostname (config-service)# icmp-object echo-reply
```


Configuring a Protocol Group

A protocol group contains IP protocol types.

Detailed Steps

	Command	Purpose
Step 1	object-group protocol <i>obj_grp_id</i> Example: ciscoasa(config)# object-group protocol tcp_udp_icmp	Adds a protocol group. The <i>obj_grp_id</i> is a text string up to 64 characters in length and can be any combination of letters, digits, and the following characters: <ul style="list-style-type: none"> • underscore “_” • dash “-” • period “.” The prompt changes to protocol configuration mode.
Step 2	Add one or more of the following group members: protocol-object <i>protocol</i> Example: ciscoasa(config-protocol)# protocol-object tcp group-object <i>group_id</i> Example: ciscoasa(config-network)# group-object Engineering_groups	Defines the protocols in the group. Enter the command for each protocol. The protocol is the numeric identifier of the specified IP protocol (1 to 254) or a keyword identifier (for example, icmp , tcp , or udp). To include all IP protocols, use the keyword ip . For a list of protocols that you can specify, see the “Protocols and Applications” section on page 49-11 . Adds an existing object group under this object group. The nested group must be of the same type.
Step 3	description <i>text</i> Example: ciscoasa(config-protocol)# description New Group	(Optional) Adds a description. The description can be up to 200 characters.

Example

To create a protocol group for TCP, UDP, and ICMP, enter the following commands:

```
hostname (config)# object-group protocol tcp_udp_icmp
hostname (config-protocol)# protocol-object tcp
hostname (config-protocol)# protocol-object udp
hostname (config-protocol)# protocol-object icmp
```

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

	Command	Purpose
Step 1	<pre>object-group user <i>user_group_name</i></pre> <p>Example: hostname(config)# object-group user users1</p>	Defines object groups that you can use to control access with the Identity Firewall.
Step 2	<p>Add one or more of the following group members:</p> <pre>user <i>domain_NetBIOS_name\user_name</i></pre> <p>Example: hostname(config-user-object-group)# user SAMPLE\users1</p>	<p>Specifies the user to add to the access rule.</p> <p>The <i>user_name</i> can contain any character including [a-z], [A-Z], [0-9], [!@#%&()-_{ } .]. If <i>domain_NetBIOS_name\user_name</i> contains a space, you must enclose the domain name and user name in quotation marks.</p> <p>The <i>user_name</i> can be part of the LOCAL domain or a user imported by the ASA from Active Directory domain.</p> <p>If the <i>domain_NetBIOS_name</i> is associated with a AAA server, the <i>user_name</i> must be the Active Directory sAMAccountName, which is unique, instead of the common name (cn), which might not be unique.</p> <p>The <i>domain_NetBIOS_name</i> can be LOCAL or the actual domain name as specified in user-identity domain <i>domain_NetBIOS_name</i> aaa-server <i>aaa_server_group_tag</i> command.</p>
	<pre>group-object <i>group_id</i></pre> <p>Example: ciscoasa(config-network)# group-object Engineering_groups</p>	Adds an existing object group under this object group. The nested group must be of the same type.
Step 3	<pre>description <i>text</i></pre> <p>Example: ciscoasa(config-protocol)# description New Group</p>	(Optional) Adds a description. The description can be up to 200 characters.

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

	Command	Purpose
Step 1	object-group security <i>objgrp_name</i> Example: ciscoasa(config)# object-group security mktg-sg	Creates a security group object. Where <i>objgrp_name</i> is the name for the group entered as a 32-byte case sensitive string. The <i>objgrp_name</i> can contain any character including [a-z], [A-Z], [0-9], [!@#%^&()-_{}].
Step 2	Add one or more of the following group members: security-group { tag <i>sgt#</i> name <i>sg_name</i> } Example: ciscoasa(config)# security-group name mktg	Specifies the type of security group object as either an inline tag or a named object. <ul style="list-style-type: none"> tag <i>sgt#</i>—Enter a number from 1 to 65533 for a Tag security type. name <i>sg_name</i>—Enter a 32-byte case-sensitive string for a Name security type. The <i>sg_name</i> can contain any character including [a-z], [A-Z], [0-9], [!@#%^&()-_{}]. 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.

Command	Purpose
group-object <i>group_id</i> Example: ciscoasa(config-network)# group-object Engineering_groups	Adds an existing object group under this object group. The nested group must be of the same type.
Step 3 description <i>text</i> Example: ciscoasa(config-protocol)# description New Group	(Optional) Adds a description. The description can be up to 200 characters.

Examples

The following example shows how to configure a security group object:

```
ciscoasa(config)# object-group security mktg-sg
ciscoasa(config)# security-group name mktg
ciscoasa(config)# security-group tag 1
```

The following example shows how to configure a security group object:

```
ciscoasa(config)# object-group security mktg-sg-all
ciscoasa(config)# security-group name mktg-managers
ciscoasa(config)# group-object mktg-sg // nested object-group
```

Configuring Regular Expressions

- [Creating a Regular Expression, page 17-14](#)
- [Creating a Regular Expression Class Map, page 17-17](#)

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

Use **Ctrl+V** to escape all of the special characters in the CLI, such as question mark (?) or a tab. For example, type **d[Ctrl+V]?g** to enter **d?g** in the configuration.

See the **regex** command in the command reference for performance impact information when matching a regular expression to packets.



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 17-1 lists the metacharacters that have special meanings.

Table 17-1 *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.
(<i>exp</i>)	Subexpression	A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, d(ola)g matches dog and dag, but dolag 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 abxyxyz.
	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.
{ <i>x</i> } or { <i>x</i> ,}	Minimum repeat quantifier	Repeat at least <i>x</i> times. For example, ab(xy){2,}z matches abxyxyz, abxyxyz, and so on.
[<i>abc</i>]	Character class	Matches any character in the brackets. For example, [abc] matches a, b, or c.
[^ <i>abc</i>]	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.
[<i>a-c</i>]	Character range class	Matches any character in the range. [a-z] matches any lowercase letter. You can mix characters and ranges: [abcq-z] matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does [a-cq-z] . The dash (-) character is literal only if it is the last or the first character within the brackets: [abc-] or [-abc] .
“”	Quotation marks	Preserves trailing or leading spaces in the string. For example, “ test” preserves the leading space when it looks for a match.
^	Caret	Specifies the beginning of a line.

Table 17-1 *regex Metacharacters (continued)*

Character	Description	Notes
\	Escape character	When used with a metacharacter, matches a literal character. For example, \[matches the left square bracket.
<i>char</i>	Character	When character is not a metacharacter, matches the literal character.
\r	Carriage return	Matches a carriage return 0x0d.
\n	Newline	Matches a new line 0x0a.
\t	Tab	Matches a tab 0x09.
\f	Formfeed	Matches a form feed 0x0c.
\xNN	Escaped hexadecimal number	Matches an ASCII character using hexadecimal (exactly two digits).
\NNN	Escaped octal number	Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space.

Detailed Steps

- Step 1** To test a regular expression to make sure it matches what you think it will match, enter the following command:

```
ciscoasa(config)# test regex input_text regular_expression
```

Where the *input_text* argument is a string you want to match using the regular expression, up to 201 characters in length.

The *regular_expression* argument can be up to 100 characters in length.

Use **Ctrl+V** to escape all of the special characters in the CLI. For example, to enter a tab in the input text in the **test regex** command, you must enter **test regex “test[Ctrl+V Tab]” “test\t”**.

If the regular expression matches the input text, you see the following message:

```
INFO: Regular expression match succeeded.
```

If the regular expression does not match the input text, you see the following message:

```
INFO: Regular expression match failed.
```

- Step 2** To add a regular expression after you tested it, enter the following command:

```
ciscoasa(config)# regex name regular_expression
```

Where the *name* argument can be up to 40 characters in length.

The *regular_expression* argument can be up to 100 characters in length.

Examples

The following example creates two regular expressions for use in an inspection policy map:

```
ciscoasa(config)# regex url_example example\.com
```

```
ciscoasa(config)# regex url_example2 example2\.com
```

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 17-14.

Detailed Steps

-
- Step 1** Create a class map by entering the following command:

```
ciscoasa(config)# class-map type regex match-any class_map_name  
ciscoasa(config-cmap)#
```

Where *class_map_name* is a string 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.

The **match-any** keyword specifies that the traffic matches the class map if it matches at least one of the regular expressions.

The CLI enters class-map configuration mode.

- Step 2** (Optional) Add a description to the class map by entering the following command:

```
ciscoasa(config-cmap)# description string
```

- Step 3** Identify the regular expressions you want to include by entering the following command for each regular expression:

```
ciscoasa(config-cmap)# match regex regex_name
```

Examples

The following example creates two regular expressions, and adds them to a regular expression class map. Traffic matches the class map if it includes the string “example.com” or “example2.com.”

```
ciscoasa(config)# regex url_example example\.com  
ciscoasa(config)# regex url_example2 example2\.com  
ciscoasa(config)# class-map type regex match-any URLs  
ciscoasa(config-cmap)# match regex url_example  
ciscoasa(config-cmap)# match regex url_example2
```

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

	Command	Purpose
Step 1	time-range <i>name</i> Example: ciscoasa(config)# time range Sales	Identifies the time-range name.
Step 2	Do one of the following: periodic <i>days-of-the-week time to</i> <i>[days-of-the-week] time</i> Example: ciscoasa(config-time-range)# periodic monday 7:59 to friday 17:01	Specifies a recurring time range. You can specify the following values for <i>days-of-the-week</i> : <ul style="list-style-type: none"> • monday, tuesday, wednesday, thursday, friday, saturday, or sunday. • daily • weekdays • weekend The <i>time</i> is in the format <i>hh:mm</i> . For example, 8:00 is 8:00 a.m. and 20:00 is 8:00 p.m.
	absolute start <i>time date [end time date]</i> Example: ciscoasa(config-time-range)# absolute start 7:59 2 january 2009	Specifies an absolute time range. The <i>time</i> is in the format <i>hh:mm</i> . For example, 8:00 is 8:00 a.m. and 20:00 is 8:00 p.m. The <i>date</i> is in the format <i>day month year</i> ; for example, 1 january 2006 .

Examples

The following is an example of an absolute time range beginning at 8:00 a.m. on January 1, 2006. Because no end time and date are specified, the time range is in effect indefinitely.


```
ciscoasa(config)# time-range for2006
ciscoasa(config-time-range)# absolute start 8:00 1 january 2006
```

The following is an example of a weekly periodic time range from 8:00 a.m. to 6:00 p.m. on weekdays:

```
ciscoasa(config)# time-range workinghours
ciscoasa(config-time-range)# periodic weekdays 8:00 to 18:00
```

Monitoring Objects

To monitor objects and groups, enter the following commands:

Command	Purpose
<code>show access-list</code>	Displays the access list entries that are expanded out into individual entries without their object groupings.
<code>show running-config object-group</code>	Displays all current object groups.
<code>show running-config object-group grp_id</code>	Displays the current object groups by their group ID.
<code>show running-config object-group grp_type</code>	Displays the current object groups by their group type.

Feature History for Objects

Table 17-2 lists each feature change and the platform release in which it was implemented.

Table 17-2 Feature History for Object Groups

Feature Name	Platform Releases	Feature Information
Object groups	7.0(1)	Object groups simplify ACL creation and maintenance. We introduced or modified the following commands: object-group protocol , object-group network , object-group service , object-group icmp_type .
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. We introduced or modified the following commands: object-network , object-service , object-group network , object-group service , network object , access-list extended , access-list webtype , access-list remark .
User Object Groups for Identity Firewall	8.4(2)	User object groups for identity firewall were introduced. We introduced the following commands: object-network user , user .

Table 17-2 Feature History for Object Groups (continued)

Feature Name	Platform Releases	Feature Information
Mixed IPv4 and IPv6 network object groups	9.0(1)	<p>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.</p> <p>Note You cannot use a mixed object group for NAT.</p> <p>We modified the following commands: object-group network.</p>
Security Group Object Groups for Cisco TrustSec	8.4(2)	<p>Security group object groups for TrustSec were introduced.</p> <p>We introduced the following commands: object-network security, security.</p>
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	<p>ICMP traffic can now be permitted/denied based on ICMP code.</p> <p>We introduced or modified the following commands: access-list extended, service-object, service.</p>



Information About Access Control Lists

Cisco ASAs provide basic traffic filtering capabilities with access control lists (ACLs), which control access in your network by preventing certain traffic from entering or exiting. This chapter describes ACLs and shows how to add them to your network configuration.

ACLs are made up of one or more access control entries (ACEs). An ACE is a single entry in an ACL that specifies a permit or deny rule (to forward or drop the packet) and is applied to a protocol, to a source and destination IP address or network, and, optionally, to the source and destination ports.

ACLs can be configured for all routed and network protocols (IP, AppleTalk, and so on) to filter the packets of those protocols as the packets pass through a router.

ACLs are used in a variety of features. If your feature uses Modular Policy Framework, you can use an ACL to identify traffic within a traffic class map. For more information on Modular Policy Framework, see Chapter 1, “Configuring a Service Policy Using the Modular Policy Framework,” in the firewall configuration guide.

This chapter includes the following sections:

- [ACL Types, page 18-1](#)
- [Access Control Entry Order, page 18-2](#)
- [Access Control Implicit Deny, page 18-3](#)
- [IP Addresses Used for ACLs When You Use NAT, page 18-3](#)
- [Where to Go Next, page 18-3](#)

ACL Types

The ASA uses five types of access control lists:

- **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. For more information, see [Chapter 21, “Adding a Standard Access Control List.”](#)
- **Extended ACLs**—Use one or more access control entries (ACE) 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). For more information, see [Chapter 19, “Adding an Extended Access Control List.”](#)
- **EtherType ACLs**—Use one or more ACEs that specify an EtherType. For more information, see [Chapter 20, “Adding an EtherType Access Control List.”](#)

- Webtype ACLs—Used in a configuration that supports filtering for clientless SSL VPN. For more information, see [Chapter 22, “Adding a Webtype Access Control List.”](#)

[Table 18-1](#) lists the types of ACLs and some common uses for them.

Table 18-1 *ACL Types and Common Uses*

ACL Use	ACL Type	Description
Control network access for IP traffic (routed and transparent mode)	Extended	The ASA does not allow any traffic from a lower security interface to a higher security interface unless it is explicitly permitted by an extended ACL. Note To access the ASA interface for management access, you do not also need an ACL allowing the host IP address. You only need to configure management access according to Chapter 41, “Configuring Management Access.”
Identify traffic for AAA rules	Extended	AAA rules use ACLs to identify traffic.
Control network access for IP traffic for a given user	Extended, downloaded from a AAA server per user	You can configure the RADIUS server to download a dynamic ACL to be applied to the user, or the server can send the name of an ACL that you already configured on the ASA.
Identify addresses for NAT (policy NAT and NAT exemption)	Extended	Policy NAT lets you identify local traffic for address translation by specifying the source and destination addresses in an extended ACL.
Establish VPN access	Extended	You can use an extended ACL in VPN commands.
Identify traffic in a traffic class map for Modular Policy Framework	Extended EtherType	ACLs can be used to identify traffic in a class map, which is used for features that support Modular Policy Framework. Features that support Modular Policy Framework include TCP and general connection settings, and inspection.
For transparent firewall mode, control network access for non-IP traffic	EtherType	You can configure an ACL that controls traffic based on its EtherType.
Identify OSPF route redistribution	Standard	Standard ACLs include only the destination address. You can use a standard ACL to control the redistribution of OSPF routes.
Filtering for WebVPN	Webtype	You can configure a Webtype ACL to filter URLs.
Control network access for IPV6 networks	IPv6	You can add and apply ACLs to control traffic in IPv6 networks.

Access Control Entry Order

An ACL is made up of one or more access control entries (ACEs). Each ACE that you enter for a given ACL name is appended to the end of the ACL. Depending on the ACL type, you can specify the source and destination addresses, the protocol, the ports (for TCP or UDP), the ICMP type (for ICMP), or the EtherType.

The order of ACEs is important. When the ASA decides whether to forward or to drop a packet, the ASA tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an ACL that explicitly permits all traffic, no further statements are checked, and the packet is forwarded.

Access Control Implicit Deny

All ACLs have an implicit deny statement at the end, so unless you explicitly permit traffic to pass, it will be denied. For example, if you want to allow all users to access a network through the ASA except for one or more particular addresses, then you need to deny those particular addresses and then permit all others.

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

IP Addresses Used for ACLs When You Use NAT

For the following features, you should always use the *real* IP address in the ACL when you use NAT, even if the address as seen on an interface is the mapped address:

- **access-group** command
- Modular Policy Framework **match access-list** command
- Botnet Traffic Filter **dynamic-filter enable classify-list** command
- AAA **aaa ... match** commands
- WCCP **wccp redirect-list group-list** command

The following features use ACLs, but these ACLs use the *mapped* values as seen on an interface:

- IPsec ACLs
- capture command ACLs
- Per-user ACLs
- Routing protocols
- All other features...

Where to Go Next

For information about implementing ACLs, see the following chapters:

- [Chapter 19, “Adding an Extended Access Control List”](#)
- [Chapter 20, “Adding an EtherType Access Control List”](#)
- [Chapter 21, “Adding a Standard Access Control List”](#)
- [Chapter 22, “Adding a Webtype Access Control List”](#)
- Chapter 6, “Configuring Access Rules,” in the firewall configuration guide



Adding an Extended Access Control List

This chapter describes how to configure extended access control lists (ACLs), and it includes the following sections:

- [Information About Extended ACLs, page 19-1](#)
- [Licensing Requirements for Extended ACLs, page 19-3](#)
- [Guidelines and Limitations, page 19-3](#)
- [Default Settings, page 19-4](#)
- [Configuring Extended ACLs, page 19-4](#)
- [Monitoring Extended ACLs, page 19-10](#)
- [Configuration Examples for Extended ACLs, page 19-10](#)
- [Where to Go Next, page 19-12](#)
- [Feature History for Extended ACLs, page 19-12](#)

Information About Extended ACLs

ACLs are used to control network access or to specify traffic for many features to act upon. An extended ACL is made up of one or more access control entries (ACEs). Each ACE specifies a source and destination for matching traffic. You can identify parameters within the **access-list** command, or you can create objects or object groups for use in the ACL.

- [Access Control Entry Order, page 19-1](#)
- [NAT and ACLs, page 19-2](#)

Access Control Entry Order

An ACL is made up of one or more ACEs. Each ACE that you enter for a given ACL name is appended to the end of the ACL.

The order of ACEs is important. When the security appliance decides whether to forward or drop a packet, the security appliance tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an ACL that explicitly permits all traffic, no further statements are ever checked.

You can disable an ACE by making it inactive.

NAT and ACLs

When using NAT or PAT, mapped addresses and ports are no longer required in an ACL for several features. You should now always use the real, untranslated addresses and ports for these features. Using the real address and port means that if the NAT configuration changes, you do not need to change the ACLs.



Note

For ACL migration information, see the *Cisco ASA 5500 Migration to Version 8.3 and Later*.

Features That Use Real IP Addresses

The following commands and features use real IP addresses in the ACLs:

- **access-group** command
- Modular Policy Framework **match access-list** command
- Botnet Traffic Filter **dynamic-filter enable classify-list** command
- AAA **aaa ... match** commands
- WCCP **wccp redirect-list group-list** command

For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5):

```
ciscoasa(config)# object network server1
ciscoasa(config-network-object)# host 10.1.1.5
ciscoasa(config-network-object)# nat (inside,outside) static 209.165.201.5

ciscoasa(config)# access-list OUTSIDE extended permit tcp any host 10.1.1.5 eq www
ciscoasa(config)# access-group OUTSIDE in interface outside
```

Features That Use Mapped IP Addresses

The following features use ACLs, but these ACLs will continue to use the mapped values as seen on an interface:

- IPsec ACLs
- **capture** command ACLs
- Per-user ACLs
- Routing protocol ACLs
- All other feature ACLs...

Information About Scheduling ACL Activation

You can schedule each ACE in an ACL to be activated at specific times of the day and week by applying a time range to the ACE.

Users could experience a delay of approximately 80 to 100 seconds after the specified end time for the ACL to become inactive. For example, if the specified end time is 3:50, because the end time is inclusive, the command is picked up anywhere between 3:51:00 and 3:51:59. After the command is picked up, the ASA finishes any currently running task and then services the command to deactivate the ACL.

Licensing Requirements for Extended ACLs

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 modes.

IPv6 Guidelines

Supports IPv6.

Features That Do Not Support IDFW, FQDN, and TrustSec ACLs

The following features use ACLs, but cannot accept an ACL with IDFW, FQDN, or TrustSec values:

- **route-map** command
- VPN **crypto map** command
- VPN **group-policy** command, except for **vpn-filter**
- WCCP
- DAP

Additional Guidelines and Limitations

- **Tip:** 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).
- Typically, you identify the **ip** keyword for the protocol, but other protocols are accepted. For a list of protocol names, see the [“Protocols and Applications” section on page 49-11](#).
- 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 49-11](#). 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.
- When you specify a network mask, the method is different from the Cisco IOS software **access-list** command. The ASA uses a network mask (for example, 255.255.255.0 for a Class C mask). The Cisco IOS mask uses wildcard bits (for example, 0.0.0.255).

Default Settings

Table 19-1 lists the default settings for extended ACL parameters.

Table 19-1 Default Extended ACL Parameters

Parameters	Default
ACE logging	ACE logging generates system log message 106023 for denied packets. A deny ACE must be present to log denied packets.
log	When the log keyword is specified, the default level for system log message 106100 is 6 (informational), and the default interval is 300 seconds.

Configuring Extended ACLs

This section shows how to add ACEs of various types to an ACL and includes the following topics:

- [Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy, page 19-4](#)
- [Adding an ACE for TCP or UDP-Based Policy, with Ports, page 19-6](#)
- [Adding an ACE for ICMP-Based Policy, with ICMP Type, page 19-7](#)
- [Adding an ACE for User-Based Policy \(Identity Firewall\), page 19-7](#)
- [Adding an ACE for Security Group-Based Policy \(TrustSec\), page 19-8](#)
- [Adding Remarks to ACLs, page 19-9](#)

Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy

This section lets you control traffic based on IP addresses or fully qualified domain names (FQDNs). An ACL is made up of one or more access control entries (ACEs) with the same ACL ID. 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 entries to the list.

Prerequisites

(Optional) Create network objects or object groups according to the [“Configuring Network Objects and Groups” section on page 17-2](#). Objects can contain an IP address (host, subnet, or range) or an FQDN. Object groups contain multiple objects or inline entries.

Guidelines

To delete an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration. To remove the entire ACL, use the **clear configure access-list** command.

Detailed Steps

Command	Purpose
<pre>access-list access_list_name [line line_number] extended {deny permit} protocol_argument source_address_argument dest_address_argument [log [[level]] [interval secs] disable default]] [inactive time-range time_range_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list ACL_IN extended permit ip any any</pre>	<p>Adds an ACE for IP address or FQDN policy.</p> <ul style="list-style-type: none"> • Line number—The line <i>line_number</i> option specifies the line number at which insert the ACE; otherwise, the ACE is added to the end of the ACL. • Permit or Deny—The deny keyword denies or exempts a packet if the conditions are matched. The permit keyword permits a packet if the conditions are matched. • Protocol—The <i>protocol_argument</i> specifies the IP protocol: <ul style="list-style-type: none"> – <i>name</i> or <i>number</i>—Specifies the protocol name or number. Specify ip to apply to all protocols. – object-group <i>protocol_grp_id</i>—Specifies a protocol object group created using the object-group protocol command. – object <i>service_obj_id</i>—Specifies a service object created using the object service command. A TCP, UDP, or ICMP service object can include a protocol <i>and</i> a source and/or destination port or ICMP type and code. – object-group <i>service_grp_id</i>—Specifies a service object group created using the object-group service command. • Source Address, Destination Address—The <i>source_address_argument</i> specifies the IP address or FQDN from which the packet is being sent, and the <i>dest_address_argument</i> specifies the IP address or FQDN to which the packet is being sent: <ul style="list-style-type: none"> – host <i>ip_address</i>—Specifies an IPv4 host address. – <i>dest_ip_address mask</i>—Specifies an IPv4 network address and subnet mask. – <i>ipv6-address/prefix-length</i>—Specifies an IPv6 host or network address and prefix. – any, any4, and any6—any specifies both IPv4 and IPv6 traffic; any4 specifies only IPv4 traffic; and any6 specifies any6 traffic. – object <i>nw_obj_id</i>—Specifies a network object created using the object network command. – object-group <i>nw_grp_id</i>—Specifies a network object group created using the object-group network command. • Logging—log arguments set logging options when an ACE matches a packet for network access (an ACL applied with the access-group command). • Activation—Inactivates or enables a time range that the ACE is active; see the time-range command for information about defining a time range.

Adding an ACE for TCP or UDP-Based Policy, with Ports

This section lets you control traffic based on IP addresses or fully qualified domain names (FQDNs) along with TCP or UDP ports. An ACL is made up of one or more access control entries (ACEs) with the same ACL ID. 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 entries to the list.

Prerequisites

- (Optional) Create network objects or object groups according to the “[Configuring Network Objects and Groups](#)” section on page 17-2. Objects can contain an IP address (host, subnet, or range) or an FQDN. Object groups contain multiple objects or inline entries.
- (Optional) Create service objects or groups according to the “[Configuring Service Objects and Service Groups](#)” section on page 17-5.

Guidelines

To delete an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration. To remove the entire ACL, use the **clear configure access-list** command.

Detailed Steps

Command	Purpose
<pre>access-list access_list_name [line line_number] extended {deny permit} {tcp udp} source_address_argument [port_argument] dest_address_argument [port_argument] [log [[level]]] [interval secs] disable default] [inactive time-range time_range_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list ACL_IN extended deny tcp any host 209.165.201.29 eq www</pre>	<p>Adds an ACE for IP address or FQDN policy, as well as optional TCP or UDP ports. For common keywords and arguments, see the “Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy” section on page 19-4. Keywords and arguments specific to this type of ACE include the following:</p> <p><i>port_argument</i> specifies the source and/or destination port. Available arguments include:</p> <ul style="list-style-type: none"> • <i>operator port</i>—The <i>operator</i> can be one of the following: <ul style="list-style-type: none"> – lt—less than – gt—greater than – eq—equal to – neq—not equal to – range—an inclusive range of values. When you use this operator, specify two port numbers, for example: <pre>range 100 200</pre> <p>The <i>port</i> can be the integer or name of a TCP or UDP port. 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.</p> <ul style="list-style-type: none"> • object-group <i>service_grp_id</i>—Specifies a service object group created using the object-group service command.

Adding an ACE for ICMP-Based Policy, with ICMP Type

This section lets you control traffic based on IP addresses or fully qualified domain names (FQDNs) along with the ICMP type. An ACL is made up of one or more access control entries (ACEs) with the same ACL ID. 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 entries to the list.

Prerequisites

- (Optional) Create network objects or object groups according to the “[Configuring Network Objects and Groups](#)” section on page 17-2. Objects can contain an IP address (host, subnet, or range) or an FQDN. Object groups contain multiple objects or inline entries.
- (Optional) Create ICMP groups according to the “[Configuring an ICMP Group](#)” section on page 17-10.

Guidelines

To delete an ACE, enter the **no access-list** command with the entire command syntax string as it appears in the configuration. To remove the entire ACL, use the **clear configure access-list** command.

Detailed Steps

Command	Purpose
<pre>access-list access_list_name [line line_number] extended {deny permit} icmp source_address_argument dest_address_argument [icmp_argument] [log [[level] [interval secs] disable default]] [inactive time-range time_range_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list abc extended permit icmp any any object-group obj_icmp_1</pre>	<p>Adds an ACE for IP address or FQDN policy, as well as optional TCP or UDP ports. For common keywords and arguments, see the “Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy” section on page 19-4. Keywords and arguments specific to this type of ACE include the following:</p> <p><i>icmp_argument</i> specifies the ICMP type and code.</p> <ul style="list-style-type: none"> • <i>icmp_type</i> [<i>icmp_code</i>]<i>—</i>Specifies the ICMP type by name or number, and the optional ICMP code for that type. If you do not specify the code, then all codes are used. • object-group <i>icmp_grp_id</i><i>—</i>Specifies an ICMP object group created using the object-group icmp command.

Adding an ACE for User-Based Policy (Identity Firewall)

If you configure the identity firewall feature, you can control traffic based on user identity.

Prerequisites

See [Chapter 38, “Configuring the Identity Firewall,”](#) to enable IDFW.

Detailed Steps

Command	Purpose
<p>access-list <i>access_list_name</i> [line <i>line_number</i>] extended {deny permit} <i>protocol_argument</i> [<i>user_argument</i>] <i>source_address_argument</i> [<i>port_argument</i>] <i>dest_address_argument</i> [<i>port_argument</i>] [log [[<i>level</i>] [interval secs] disable default]] [inactive time-range <i>time_range_name</i>]</p> <p>Example:</p> <pre>ciscoasa(config)# access-list v1 extended permit ip user LOCAL\idfw any 10.0.0.0 255.255.255.0</pre>	<p>Adds an ACE for IP address or FQDN policy, as well as optional usernames and/or groups. For common keywords and arguments, see the “Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy” section on page 19-4. Keywords and arguments specific to this type of ACE include the following:</p> <p><i>user_argument</i> is for use with the identity firewall feature, and specifies the user or group for which to match traffic in addition to the source address. Available arguments include the following:</p> <ul style="list-style-type: none"> • object-group-user <i>user_obj_grp_id</i>—Specifies a user object group created using the object-group user command. • user {[<i>domain_nickname</i>\]<i>name</i> any none}—Specifies a username. Specify any to match all users with user credentials, or none to match users without user credentials. These options are especially useful for combining access-group and aaa authentication match policies. • user-group [<i>domain_nickname</i>\]<i>user_group_name</i>—Specifies a user group name. <p>Note Although not shown in the syntax at left, you can also use TrustSec security group arguments.</p>

Adding an ACE for Security Group-Based Policy (TrustSec)

If you configure the Cisco TrustSec feature, you can control traffic based on security groups.

Prerequisites

See [Chapter 39, “Configuring the ASA to Integrate with Cisco TrustSec,”](#) to enable TrustSec.

Detailed Steps

Command	Purpose
<pre>access-list access_list_name [line line_number] extended {deny permit} protocol_argument [security_group_argument] source_address_argument [port_argument] [security_group_argument] dest_address_argument [port_argument] [log [[level] [interval secs] disable default]] [inactive time-range time_range_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list v1 extended permit ip user LOCAL\idfw any 10.0.0.0 255.255.255.0</pre>	<p>Adds an ACE for IP address or FQDN policy, as well as optional security groups. For common keywords and arguments, see the “Adding an ACE for IP Address or Fully Qualified Domain Name-Based Policy” section on page 19-4. Keywords and arguments specific to this type of ACE include the following:</p> <p><i>security_group_argument</i> is for use with the TrustSec feature, and specifies the security group for which to match traffic in addition to the source or destination address. Available arguments include the following:</p> <ul style="list-style-type: none"> • object-group-security <i>security_obj_grp_id</i>—Specifies a security object group created using the object-group security command. • security-group { <i>name security_grp_id</i> <i>tag security_grp_tag</i> }—Specifies a security group name or tag. <p>Note Although not shown in the syntax at left, you can also use Identity Firewall user arguments.</p>

Adding Remarks to ACLs

You can include remarks about entries in any ACL. The remarks make the ACL easier to understand.

To add a remark after the last **access-list** command you entered, enter the following command.

Detailed Steps

Command	Purpose
<pre>access-list access_list_name remark text</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list OUT remark - this is the inside admin address</pre>	<p>Adds a remark after the last access-list command you entered.</p> <p>The text can be up to 100 characters in length. You can enter leading spaces at the beginning of the text. Trailing spaces are ignored.</p> <p>If you enter the remark before any access-list command, then the remark is the first line in the ACL.</p> <p>If you delete an ACL using the no access-list <i>access_list_name</i> command, then all the remarks are also removed.</p>

Examples

You can add remarks before each ACE, and the remark appears in the ACL in this location. Entering a dash (-) at the beginning of the remark helps set it apart from the ACEs.

```
ciscoasa(config)# access-list OUT remark - this is the inside admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.3 any
ciscoasa(config)# access-list OUT remark - this is the hr admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.4 any
```

Monitoring Extended ACLs

To monitor extended ACLs, enter one of the following commands:

Command	Purpose
<code>show access-list</code>	Displays the ACEs by number.
<code>show running-config access-list</code>	Displays the current running access-list configuration.

Configuration Examples for Extended ACLs

This section includes the following topics:

- [Configuration Examples for Extended ACLs \(No Objects\), page 19-10](#)
- [Configuration Examples for Extended ACLs \(Using Objects\), page 19-11](#)

Configuration Examples for Extended ACLs (No Objects)

The following ACL allows all hosts (on the interface to which you apply the ACL) to go through the ASA:

```
hostname(config)# access-list ACL_IN extended permit ip any any
```

The following sample ACL prevents hosts on 192.168.1.0/24 from accessing the 209.165.201.0/27 network. All other addresses are permitted.

```
hostname(config)# access-list ACL_IN extended deny tcp 192.168.1.0 255.255.255.0
209.165.201.0 255.255.255.224
hostname(config)# access-list ACL_IN extended permit ip any any
```

If you want to restrict access to selected hosts only, then enter a limited permit ACE. By default, all other traffic is denied unless explicitly permitted.

```
hostname(config)# access-list ACL_IN extended permit ip 192.168.1.0 255.255.255.0
209.165.201.0 255.255.255.224
```

The following ACL restricts all hosts (on the interface to which you apply the ACL) from accessing a website at address 209.165.201.29. All other traffic is allowed.

```
hostname(config)# access-list ACL_IN extended deny tcp any host 209.165.201.29 eq www
hostname(config)# access-list ACL_IN extended permit ip any any
```

The following ACL that uses object groups restricts several hosts on the inside network from accessing several web servers. All other traffic is allowed.

```
hostname(config-network)# access-list ACL_IN extended deny tcp object-group denied
object-group web eq www
hostname(config)# access-list ACL_IN extended permit ip any any
hostname(config)# access-group ACL_IN in interface inside
```

The following example temporarily disables an ACL that permits traffic from one group of network objects (A) to another group of network objects (B):

```
hostname(config)# access-list 104 permit ip host object-group A object-group B inactive
```


To implement a time-based ACL, use the **time-range** command to define specific times of the day and week. Then use the **access-list extended** command to bind the time range to an ACL. The following example binds an ACL named “Sales” to a time range named “New_York_Minute.”

```
hostname(config)# access-list Sales line 1 extended deny tcp host 209.165.200.225 host
209.165.201.1 time-range New_York_Minute
```

The following example shows a mixed IPv4/IPv6 ACL:

```
hostname(config)# access-list demoacl extended permit ip 2001:DB8:1::/64 10.2.2.0
255.255.255.0
hostname(config)# access-list demoacl extended permit ip 2001:DB8:1::/64 2001:DB8:2::/64
hostname(config)# access-list demoacl extended permit ip host 10.3.3.3 host 10.4.4.4
```

Configuration Examples for Extended ACLs (Using Objects)

The following normal ACL that does not use object groups restricts several hosts on the inside network from accessing several web servers. All other traffic is allowed.

```
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.29
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.29
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.29
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.16
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.16
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.16
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.4 host 209.165.201.78
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.78 host 209.165.201.78
eq www
ciscoasa(config)# access-list ACL_IN extended deny tcp host 10.1.1.89 host 209.165.201.78
eq www
ciscoasa(config)# access-list ACL_IN extended permit ip any any
ciscoasa(config)# access-group ACL_IN in interface inside
```

If you make two network object groups, one for the inside hosts, and one for the web servers, then the configuration can be simplified and can be easily modified to add more hosts:

```
ciscoasa(config)# object-group network denied
ciscoasa(config-network)# network-object host 10.1.1.4
ciscoasa(config-network)# network-object host 10.1.1.78
ciscoasa(config-network)# network-object host 10.1.1.89

ciscoasa(config-network)# object-group network web
ciscoasa(config-network)# network-object host 209.165.201.29
ciscoasa(config-network)# network-object host 209.165.201.16
ciscoasa(config-network)# network-object host 209.165.201.78

ciscoasa(config-network)# access-list ACL_IN extended deny tcp object-group denied
object-group web eq www
ciscoasa(config)# access-list ACL_IN extended permit ip any any
ciscoasa(config)# access-group ACL_IN in interface inside
```

Where to Go Next

Many features use ACLs. Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide for more information.

Feature History for Extended ACLs

Table 19-2 lists the release history for this feature.

Table 19-2 Feature History for Extended ACLs

Feature Name	Releases	Feature Information
Extended ACLs	7.0(1)	<p>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 (ACE) 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 (for ICMP).</p> <p>We introduced the following command: access-list extended.</p>
Real IP addresses	8.3(1)	<p>When using NAT or PAT, mapped addresses and ports are no longer required in an ACL for several features. You should now always use the real, untranslated addresses and ports for these features. Using the real address and port means that if the NAT configuration changes, you do not need to change the ACLs. See the “Features That Use Real IP Addresses” section on page 19-2 for more information.</p>
Support for Identity Firewall	8.4(2)	<p>You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication.</p> <p>We modified the following commands: access-list extended.</p>
Support for TrustSec	9.0(1)	<p>You can now use TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules.</p> <p>We modified the following commands: access-list extended.</p>

Table 19-2 Feature History for Extended ACLs (continued)

Feature Name	Releases	Feature Information
Unified ACL for IPv4 and IPv6	9.0(1)	<p>ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The any keyword was changed to represent IPv4 and IPv6 traffic. The any4 and any6 keywords were added to represent IPv4-only and IPv6-only traffic, respectively. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration.</p> <p>We modified the following commands: access-list extended, access-list webtype.</p> <p>We removed the following commands: ipv6 access-list, ipv6 access-list webtype, ipv6-vpn-filter</p>
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	<p>ICMP traffic can now be permitted/denied based on ICMP code.</p> <p>We introduced or modified the following commands: access-list extended, service-object, service.</p>



Adding an EtherType Access Control List

This chapter describes how to configure EtherType ACLs and includes the following sections:

- [Information About EtherType ACLs, page 20-1](#)
- [Licensing Requirements for EtherType ACLs, page 20-1](#)
- [Guidelines and Limitations, page 20-2](#)
- [Default Settings, page 20-2](#)
- [Configuring EtherType ACLs, page 20-2](#)
- [Monitoring EtherType ACLs, page 20-4](#)
- [What to Do Next, page 20-4](#)
- [Configuration Examples for EtherType ACLs, page 20-5](#)
- [Feature History for EtherType ACLs, page 20-5](#)

Information About EtherType ACLs

An EtherType ACL is made up of one or more Access Control Entries (ACEs) that specify an EtherType. An EtherType rule controls any EtherType identified by a 16-bit hexadecimal number, as well as selected traffic types. See the Supported EtherTypes and Other Traffic section in the firewall configuration guide for more information.

For information about creating an access rule with the EtherType ACL, see Chapter 6, “Configuring Access Rules,” in the firewall configuration guide.

Licensing Requirements for EtherType 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

Available in single and multiple context modes.

Firewall Mode Guidelines

Supported in transparent firewall mode only.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

The following guidelines and limitations apply to EtherType ACLs:

- For EtherType ACLs, the implicit deny at the end of the ACL does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the ACL does not now block any IP traffic that you previously allowed with an extended ACL (or implicitly allowed from a high security interface to a low security interface). However, if you *explicitly* deny all traffic with an EtherType ACE, then IP and ARP traffic is denied.
- 802.3-formatted frames are not handled by the ACL because they use a length field as opposed to a type field.
- See the Supported EtherTypes and Other Traffic section in the firewall configuration guide for more information about supported traffic.

Default Settings

ACL logging generates system log message 106023 for denied packets. Deny packets must be present to log denied packets.

When you configure logging for the ACL, the default severity level for system log message 106100 is 6 (informational).

Configuring EtherType ACLs

This section includes the following topics:

- [Task Flow for Configuring EtherType ACLs, page 20-2](#)
- [Adding EtherType ACLs, page 20-3](#)
- [Adding Remarks to ACLs, page 20-4](#)

Task Flow for Configuring EtherType ACLs


Use the following guidelines to create and implement an ACL:

- Step 1** Create an ACL by adding an ACE and applying an ACL name, as shown in the “Adding EtherType ACLs” section on page 20-3.
- Step 2** Apply the ACL to an interface. (See the Configuring Access Rules section in the firewall configuration guide for more information.)

Adding EtherType ACLs

To configure an ACL that controls traffic based upon its EtherType, perform the following steps:

Detailed Steps

Command	Purpose
<pre>access-list access_list_name ethertype {deny permit} {ipx bpdu mpls-unicast mpls-multicast is-is any hex_number}</pre> <p>Example: ciscoasa(config)# ciscoasa(config)# access-list ETHER ethertype permit ipx</p>	<p>Adds an EtherType ACE.</p> <p>The <i>access_list_name</i> argument lists the name or number of an ACL. When you specify an ACL name, the ACE is added to the end of the ACL. Enter the <i>access_list_name</i> in upper case 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 for the purpose (for example, MPLS or PIX).</p> <p>The permit keyword permits access if the conditions are matched. deny denies access.</p> <p>The ipx keyword specifies access to IPX.</p> <p>The bpdu keyword specifies access to bridge protocol data units, which are allowed by default.</p> <p>The deny keyword denies access if the conditions are matched. If an EtherType ACL is configured to deny all, all ethernet frames are discarded. Only physical protocol traffic, such as auto-negotiation, is still allowed.</p> <p>The mpls-multicast keyword specifies access to MPLS multicast.</p> <p>The mpls-unicast keyword specifies access to MPLS unicast.</p> <p>The is-is keyword specifies access to IS-IS traffic.</p> <p>The any keyword specifies access to any traffic.</p> <p>The <i>hex_number</i> argument indicates any EtherType that can be identified by a 16-bit hexadecimal number greater than or equal to 0x600. (See RFC 1700, “Assigned Numbers,” at http://www.ietf.org/rfc/rfc1700.txt for a list of EtherTypes.)</p> <p> Note To remove an EtherType ACE, enter the no access-list command with the entire command syntax string as it appears in the configuration.</p>

Example

The following sample ACL allows common traffic originating on the inside interface:

```
ciscoasa(config)# access-list ETHER ethertype permit ipx
ciscoasa(config)# access-list ETHER ethertype permit mpls-unicast
ciscoasa(config)# access-group ETHER in interface inside
```

Adding Remarks to ACLs

You can include remarks about entries in any ACL, including extended, EtherType, IPv6, standard, and Webtype ACLs. The remarks make an ACL easier to understand.

To add a remark after the last **access-list** command you entered, enter the following command:

Command	Purpose
<code>access-list access_list_name remark text</code>	Adds a remark after the last access-list command you entered.
Example: <code>ciscoasa(config)# access-list OUT remark - this is the inside admin address</code>	The text can be up to 100 characters in length. You can enter leading spaces at the beginning of the text. Trailing spaces are ignored. If you enter the remark before any access-list command, then the remark is the first line in the ACL. If you delete an ACL using the no access-list access_list_name command, then all remarks are also removed.

Example

You can add remarks before each ACE, and the remarks appear in the ACL in these locations. Entering a dash (-) at the beginning of a remark helps to set it apart from the ACE.

```
ciscoasa(config)# access-list OUT remark - this is the inside admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.3 any
ciscoasa(config)# access-list OUT remark - this is the hr admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.4 any
```

What to Do Next

Apply the ACL to an interface. (See the Configuring Access Rules section in the firewall configuration guide for more information.)

Monitoring EtherType ACLs

To monitor EtherType ACLs, enter one of the following commands:

Command	Purpose
<code>show access-list</code>	Displays the ACL entries by number.
<code>show running-config access-list</code>	Displays the current running access-list configuration.

Configuration Examples for EtherType ACLs

The following example shows how to configure EtherType ACLs:

The following ACL allows some EtherTypes through the ASA, but it denies IPX:

```
ciscoasa(config)# access-list ETHER ethertype deny ipx
ciscoasa(config)# access-list ETHER ethertype permit 0x1234
ciscoasa(config)# access-list ETHER ethertype permit mpls-unicast
ciscoasa(config)# access-group ETHER in interface inside
ciscoasa(config)# access-group ETHER in interface outside
```

The following ACL denies traffic with EtherType 0x1256, but it allows all others on both interfaces:

```
ciscoasa(config)# access-list nonIP ethertype deny 1256
ciscoasa(config)# access-list nonIP ethertype permit any
ciscoasa(config)# access-group ETHER in interface inside
ciscoasa(config)# access-group ETHER in interface outside
```

Feature History for EtherType ACLs

Table 20-1 lists the release history for this feature.

Table 20-1 Feature History for EtherType ACLs

Feature Name	Releases	Feature Information
EtherType ACLs	7.0(1)	EtherType ACLs control traffic based upon its EtherType. We introduced the feature and the following command: access-list ethertype.
EtherType ACL support for IS-IS traffic	8.4(5), 9.1(2)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following command: access-list ethertype {permit deny} is-is.



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 21-1](#)
- [Licensing Requirements for Standard ACLs, page 21-1](#)
- [Guidelines and Limitations, page 21-1](#)
- [Default Settings, page 21-2](#)
- [Adding Standard ACLs, page 21-3](#)
- [What to Do Next, page 21-4](#)
- [Monitoring ACLs, page 21-4](#)
- [Configuration Examples for Standard ACLs, page 21-4](#)
- [Feature History for Standard ACLs, page 21-5](#)

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 21-2](#)
- [Firewall Mode Guidelines, page 21-2](#)

- [IPv6 Guidelines, page 21-2](#)
- [Additional Guidelines and Limitations, page 21-2](#)

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.
- To add additional ACEs at the end of the ACL, enter another **access-list** command, specifying the same ACL name.
- When used with the **access-group** command, the **deny** keyword does not allow a packet to traverse the ASA. By default, the ASA denies all packets on the originating interface unless you specifically permit access.
- When specifying a source, local, or destination address, use the following guidelines:
 - Use a 32-bit quantity in four-part, dotted-decimal format.
 - Use the keyword **any** as an abbreviation for an address and mask of 0.0.0.0.0.0.0.
 - Use the **host ip_address** option as an abbreviation for a mask of 255.255.255.255.
- You can disable an ACE by specifying the keyword **inactive** in the **access-list** command.

Default Settings

[Table 21-1](#) lists the default settings for standard ACL parameters.

Table 21-1 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:

- [Task Flow for Configuring Extended ACLs, page 21-3](#)
- [Adding a Standard ACL, page 21-3](#)
- [Adding Remarks to ACLs, page 21-4](#)

Task Flow for Configuring Extended ACLs

Use the following guidelines to create and implement an ACL:

- Create an ACL by adding an ACE and applying an ACL name. See in the [“Adding Standard ACLs” section on page 21-3](#).
- Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide for more information.

Adding a Standard ACL

To add an ACL to identify the destination IP addresses of OSPF routes, which can be used in a route map for OSPF redistribution, enter the following command:

Command	Purpose
<pre>hostname(config)# access-list access_list_name standard {deny permit} {any4 ip_address mask}</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list OSPF standard permit 192.168.1.0 255.255.255.0</pre>	<p>Adds a standard access list entry. To add another ACE to the end of the ACL, enter another access-list command, specifying the same ACL name.</p> <p>The <i>access_list_name</i> argument specifies the name of number of an ACL.</p> <p>The any4 keyword specifies access to anyone.</p> <p>The deny keyword denies access if the conditions are matched.</p> <p>The host ip_address syntax specifies access to a host IP address.</p> <p>The <i>ip_address ip_mask</i> argument specifies access to a specific IP address and subnet mask.</p> <p>The line line-num option specifies the line number at which to insert an ACE.</p> <p>The permit keyword permits access if the conditions are matched.</p> <p>To remove an ACE, enter the no access-list command with the entire command syntax string as it appears in the configuration.</p>

Adding Remarks to ACLs

You can include remarks about entries in any ACL, including extended, EtherType, IPv6, standard, and Webtype ACLs. The remarks make the ACL easier to understand.

To add a remark after the last **access-list** command you entered, enter the following command:

Command	Purpose
access-list <i>access_list_name</i> remark <i>text</i>	Adds a remark after the last access-list command you entered.
Example: ciscoasa(config)# access-list OUT remark - this is the inside admin address	The text can be up to 100 characters in length. You can enter leading spaces at the beginning of the text. Trailing spaces are ignored. If you enter the remark before any access-list command, then the remark is the first line in the ACL. If you delete an ACL using the no access-list <i>access_list_name</i> command, then all the remarks are also removed.

Example

You can add a remark before each ACE, and the remarks appear in the ACLs in these location. Entering a dash (-) at the beginning of a remark helps to set it apart from an ACE.

```
ciscoasa(config)# access-list OUT remark - this is the inside admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.3 any
ciscoasa(config)# access-list OUT remark - this is the hr admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.4 any
```

What to Do Next

Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide for more information.

Monitoring ACLs

To monitor ACLs, perform one of the following tasks:

Command	Purpose
show access-list	Displays the ACL entries by number.
show running-config access-list	Displays the current running access-list configuration.

Configuration Examples for Standard ACLs

The following example shows how to deny IP traffic through the ASA:

```
ciscoasa(config)# access-list 77 standard deny
```

The following example shows how to permit IP traffic through the ASA if conditions are matched:

```
ciscoasa(config)# access-list 77 standard permit
```

The following example shows how to specify a destination address:

```
ciscoasa(config)# access-list 77 standard permit host 10.1.10.123
```

Feature History for Standard ACLs

Table 21-2 lists the release history for this feature.

Table 21-2 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. We introduced the feature and the following command: access-list standard.



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 ACL to the configuration that supports filtering for WebVPN.

This chapter includes the following sections:

- [Licensing Requirements for Webtype ACLs, page 22-1](#)
- [Guidelines and Limitations, page 22-1](#)
- [Default Settings, page 22-3](#)
- [Using Webtype ACLs, page 22-3](#)
- [What to Do Next, page 22-6](#)
- [Monitoring Webtype ACLs, page 22-6](#)
- [Configuration Examples for Webtype ACLs, page 22-6](#)
- [Feature History for Webtype ACLs, page 22-8](#)

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 22-1](#)
- [Firewall Mode Guidelines, page 22-2](#)
- [Additional Guidelines and Limitations, page 22-2](#)

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.

**Note**

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. 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>/*
```
- The **access-list webtype** command is used to configure clientless SSL VPN filtering. The URL specified may be full or partial (no file specified), may include wildcards for the server, or may specify a port. See the “[Adding Webtype ACLs with a URL String](#)” section on page 22-4 for information about using wildcard characters in the URL string.
- Valid protocol identifiers are http, https, cifs, ica, imap4, pop3, and smtp. The RL may also contain the keyword **any** to refer to any URL. An asterisk may be used to refer to a subcomponent of a DNS name.
- Dynamic ACLs have been extended to support IPv6 ACLs. If you configure both an IPv4 ACL and an IPv6 ACL, they are converted to dynamic ACLs.
- If you use the Access Control Server (ACS), you must configure IPv6 ACLs using the cisco-av-pair attribute; downloadable ACLs are not supported in the ACS GUI.

- 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 22-1 lists the default settings for Webtype ACLs parameters.

Table 22-1 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:

- [Task Flow for Configuring Webtype ACLs, page 22-3](#)
- [Adding Webtype ACLs with a URL String, page 22-4](#)
- [Adding Webtype ACLs with an IP Address, page 22-5](#)
- [Adding Remarks to ACLs, page 22-5](#)

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 22-3.
- Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide for more information.

Adding Webtype ACLs with a URL String

To add an ACL to the configuration that supports filtering for clientless SSL VPN, enter the following command:

Command	Purpose
<pre>access-list access_list_name webtype {deny permit} url {url_string any} [log[[disable default] level] interval secs][time_range name]]</pre> <p>Example: ciscoasa(config)# access-list acl_company webtype deny url http://*.cisco.example</p>	<p>Adds an ACL to the configuration that supports filtering for clientless SSL VPN.</p> <p>The <i>url_string</i> option specifies the URL to be filtered. You can use the following wildcard characters to define more than one wildcard in the Webtype ACE:</p> <ul style="list-style-type: none"> • Enter an asterisk “*” to match no characters or any number of characters. • Enter a question mark “?” to match any one character exactly. • Enter square brackets “[]” to create a range operator that matches any one character in a range. <p>Note To match any HTTP URL, you must enter http://*/* instead of the former method of entering http://*.</p> <ul style="list-style-type: none"> • The any keyword specifies all URLs. <p>The interval option specifies the time interval at which to generate system log message 106100; valid values are from 1 to 600 seconds.</p> <p>The log [[disable default] <i>level</i>] option specifies that system log message 106100 is generated for the ACE. When the log optional keyword is specified, the default level for system log message 106100 is 6 (informational). See the log command for more information.</p> <p>The time_range name option specifies a keyword for attaching the time-range option to this access list entry.</p> <p>To remove an ACL, use the no form of this command with the complete syntax string as it appears in the configuration.</p>

Adding Webtype ACLs with an IP Address

To add an ACL to the configuration that supports filtering for clientless SSL VPN, enter the following command:

Command	Purpose
<pre>access-list access_list_name webtype {deny permit} tcp [dest_address_argument] [eq port] [log[[disable default] level] interval secs] [time_range name]]</pre> <p>Example: ciscoasa(config)# access-list acl_company webtype permit tcp any</p>	<p>Adds an ACL to the configuration that supports filtering for WebVPN.</p> <p>The <i>dest_address_argument</i> specifies the IP address to which the packet is being sent:</p> <ul style="list-style-type: none"> host ip_address—Specifies an IPv4 host address. dest_ip_address mask—Specifies an IPv4 network address and subnet mask. ipv6-address/prefix-length—Specifies an IPv6 host or network address and prefix. any, any4, and any6—any specifies both IPv4 and IPv6 traffic; any4 specifies only IPv4 traffic; and any6 specifies only IPv6 traffic. <p>The eq port can be one of the following ports: http, https, cifs, imap4, pop3, and smtp.</p> <p>The interval option specifies the time interval at which to generate system log message 106100; valid values are from 1 to 600 seconds.</p> <p>The log [[disable default] level] option specifies that system log message 106100 is generated for the ACE. The default level is 6 (informational).</p> <p>The time_range name option specifies a keyword for attaching the time-range option to this access list entry.</p> <p>To remove an ACL, use the no form of this command with the complete syntax string as it appears in the configuration.</p>

Adding Remarks to ACLs

You can include remarks about entries in any ACL, including extended, EtherType, IPv6, standard, and Webtype ACLs. The remarks make the ACL easier to understand.

To add a remark after the last **access-list** command you entered, enter the following command:

Command	Purpose
<pre>access-list access_list_name remark text</pre> <p>Example: ciscoasa(config)# access-list OUT remark - this is the inside admin address</p>	<p>Adds a remark after the last access-list command you entered.</p> <p>The text can be up to 100 characters in length. You can enter leading spaces at the beginning of the text. Trailing spaces are ignored.</p> <p>If you enter the remark before any access-list command, then the remark is the first line in the ACL.</p> <p>If you delete an ACL using the no access-list access_list_name command, then all the remarks are also removed.</p>

Example

You can add a remark before each ACE, and the remarks appear in the ACL in these locations. Entering a dash (-) at the beginning of a remark helps set it apart from an ACE.

```
ciscoasa(config)# access-list OUT remark - this is the inside admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.3 any
ciscoasa(config)# access-list OUT remark - this is the hr admin address
ciscoasa(config)# access-list OUT extended permit ip host 209.168.200.4 any
```

What to Do Next

Apply the ACL to an interface. See the Configuring Access Rules section in the firewall configuration guide for more information.

Monitoring Webtype ACLs

To monitor webtype ACLs, enter the following command:

Command	Purpose
<code>show running-config access-list</code>	Displays the access-list configuration running on the ASA.
<code>debug webvpn url</code>	Debug webtype ACL related issues.

Configuration Examples for Webtype ACLs

The following example shows how to deny access to a specific company URL:

```
ciscoasa(config)# access-list acl_company webtype deny url http://*.example.com
```

The following example shows how to deny access to a specific file:

```
ciscoasa(config)# access-list acl_file webtype deny url
https://www.example.com/dir/file.html
```

The following example shows how to deny HTTP access to any URL through port 8080:

```
ciscoasa(config)# access-list acl_company webtype deny url http://my-server:8080/*
```

The following examples show how to use wildcards in Webtype ACLs.

- The following example matches URLs such as `http://www.example.com/layouts/1033:`

```
access-list VPN-Group webtype permit url http://www.example.com/*
```
- The following example matches URLs such as `http://www.example.com/` and `http://www.example.net/`:

```
access-list test webtype permit url http://www.**ample.com/
```
- The following example matches URLs such as `http://www.cisco.com` and `ftp://wwz.example.com:`

```
access-list test webtype permit url *://ww?.c*co*/
```

- The following example matches URLs such as `http://www.cisco.com:80` and `https://www.cisco.com:81`:

```
access-list test webytype permit url *://ww?.c*co*:8[01]/
```

The range operator “[]” in the preceding example specifies that either character **0** or **1** can occur.

- The following example matches URLs such as `http://www.example.com` and `http://www.example.net`:

```
access-list test webytype permit url http://www.\[a-z\]ample?\*/
```

The range operator “[]” in the preceding example specifies that any character in the range from **a** to **z** can occur.

- The following example matches URLs such as `http://www.cisco.com/anything/crazy/url/ddtscgiz`:

```
access-list test webytype permit url htt*://*/.*cgi?*
```

**Note**

To match any http URL, you must enter **http://*/.*** instead of the former method of entering `http://*`.

The following example shows how to enforce a webytype ACL to disable access to specific CIFS shares.

In this scenario we have a root folder named “shares” that contains two sub-folders named “Marketing_Reports” and “Sales_Reports.” We want to specifically deny access to the “shares/Marketing_Reports” folder.

```
access-list CIFS_Avoid webytype deny url cifs://172.16.10.40/shares/Marketing_Reports.
```

However, due to the implicit “deny all,” the above ACL makes all of the sub-folders inaccessible (“shares/Sales_Reports” and “shares/Marketing_Reports”), including the root folder (“shares”).

To fix the problem, add a new ACL to allow access to the root folder and the remaining sub-folders:

```
access-list CIFS_Allow webytype permit url cifs://172.16.10.40/shares*
```

Feature History for Webtype ACLs

Table 22-2 lists the release history for this feature.

Table 22-2 Feature History for Webtype ACLs

Feature Name	Releases	Feature Information
Webtype ACLs	7.0(1)	<p>Webtype ACLs are ACLs that are added to a configuration that supports filtering for clientless SSL VPN.</p> <p>We introduced the feature and the following command: access-list webtype.</p>
Unified ACL for IPv4 and IPv6	9.0(1)	<p>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.</p> <p>We modified the following commands: access-list extended, access-list webtype.</p> <p>We removed the following commands: ipv6 access-list, ipv6 access-list webtype, ipv6-vpn-filter</p>
Webtype ACL enhancements	9.1(5)	<ul style="list-style-type: none"> 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. <p>Note A duplicate ACE refers to ACEs with URLs that are equivalent after normalization.</p> <p>We did not modify any commands.</p>



Configuring Logging for Access Control Lists

This chapter describes how to configure ACL logging for extended ACLs and Webtype ACLs, and it describes how to manage deny flows.

This chapter includes the following sections:

- [Configuring Logging for ACLs, page 23-1](#)
- [Managing Deny Flows, page 23-5](#)

Configuring Logging for ACLs

This section includes the following topics:

- [Information About Logging ACL Activity, page 23-1](#)
- [Licensing Requirements for ACL Logging, page 23-2](#)
- [Guidelines and Limitations, page 23-2](#)
- [Default Settings, page 23-3](#)
- [Configuring ACL Logging, page 23-3](#)
- [Monitoring ACLs, page 23-4](#)
- [Configuration Examples for ACL Logging, page 23-4](#)
- [Feature History for ACL Logging, page 23-5](#)

Information About Logging ACL Activity

By default, when traffic is denied by an extended ACE or a Webtype ACE, the ASA generates syslog message 106023 for each denied packet in the following form:

```
%ASA|PIX-4-106023: Deny protocol src [interface_name:source_address/source_port] dst  
interface_name:dest_address/dest_port [type {string}, code {code}] by access_group acl_id
```

If the ASA is attacked, the number of syslog messages for denied packets can be very large. We recommend that you instead enable logging using syslog message 106100, which provides statistics for each ACE and enables you to limit the number of syslog messages produced. Alternatively, you can disable all logging.

**Note**

Only ACEs in the ACL generate logging messages; the implicit deny at the end of the ACL does not generate a message. If you want all denied traffic to generate messages, add the implicit ACE manually to the end of the ACL, as shown in the following example:

```
ciscoasa(config)# access-list TEST deny ip any any log
```

The **log** options at the end of the extended **access-list** command enable you to set the following behavior:

- Enable message 106100 instead of message 106023
- Disable all logging
- Return to the default logging using message 106023

Syslog message 106100 uses the following form:

```
%ASA|PIX-n-106100: access-list acl_id {permitted | denied} protocol
interface_name/source_address(source_port) -> interface_name/dest_address(dest_port)
hit-cnt number ({first hit | number-second interval})
```

When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA generates a syslog message at the first hit and at the end of each interval, identifying the total number of hits during the interval and the timestamp for the last hit. At the end of each interval, the ASA resets the hit count to 0. If no packets match the ACE during an interval, the ASA deletes the flow entry.

A flow is defined by the source and destination IP addresses, protocols, and ports. Because the source port might differ for a new connection between the same two hosts, you might not see the same flow increment because a new flow was created for the connection. See the [“Managing Deny Flows” section on page 23-5](#) to limit the number of logging flows.

Permitted packets that belong to established connections do not need to be checked against ACLs; only the initial packet is logged and included in the hit count. For connectionless protocols, such as ICMP, all packets are logged, even if they are permitted, and all denied packets are logged.

See the *syslog messages guide* for detailed information about this syslog message.

Licensing Requirements for ACL Logging

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 and transparent firewall modes.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

ACE logging generates syslog message 106023 for denied packets. A deny ACE must be present to log denied packets.

Default Settings

Table 23-1 lists the default settings for extended ACL parameters.

Table 23-1 Default Extended ACL Parameters

Parameters	Default
log	When the log keyword is specified, the default level for syslog message 106100 is 6 (informational), and the default interval is 300 seconds.

Configuring ACL Logging

This sections describes how to configure ACL logging.

**Note**

For complete ACL command syntax, see the [“Configuring Extended ACLs”](#) section on page 19-4 and the [“Using Webtype ACLs”](#) section on page 22-3.

To configure logging for an ACE, enter the following command:

Command	Purpose
<pre>access-list access_list_name [extended] {deny / permit}...[log [[level] [interval secs] disable default]]</pre> <p>Example:</p> <pre>ciscoasa(config)# access-list outside-acl permit ip host 1.1.1.1 any log 7 interval 600</pre>	<p>Configures logging for an ACE.</p> <p>The access-list <i>access_list_name</i> syntax specifies the ACL for which you want to configure logging.</p> <p>The extended option adds an ACE.</p> <p>The deny keyword denies a packet if the conditions are matched. Some features do not allow deny ACEs, such as NAT. (See the command documentation for each feature that uses an ACL for more information.)</p> <p>The permit keyword permits a packet if the conditions are matched.</p> <p>If you enter the log option without any arguments, you enable syslog message 106100 at the default level (6) and for the default interval (300 seconds). See the following options:</p> <ul style="list-style-type: none"> • level—A severity level between 0 and 7. The default is 6. • interval secs—The time interval in seconds between syslog messages, from 1 to 600. The default is 300. This value is also used as the timeout value for deleting an inactive flow. • disable—Disables all ACL logging. • default—Enables logging to message 106023. This setting is the same as having no log option. <p>(See the access-list command in the <i>Cisco Security Appliance Command Reference</i> for more information about command options.)</p>

Monitoring ACLs

To monitor ACLs, enter one of the following commands:

Command	Purpose
<code>show access-list</code>	Displays the ACL entries by number.
<code>show running-config access-list</code>	Displays the current running ACL configuration.

Configuration Examples for ACL Logging

This section includes sample configurations for logging ACLs.

You might configure the following ACL:

```
ciscoasa(config)# access-list outside-acl permit ip host 1.1.1.1 any log 7 interval 600
ciscoasa(config)# access-list outside-acl permit ip host 2.2.2.2 any
ciscoasa(config)# access-list outside-acl deny ip any any log 2
ciscoasa(config)# access-group outside-acl in interface outside
```

When the first ACE of `outside-acl` permits a packet, the ASA generates the following syslog message:

```
%ASA|PIX-7-106100: access-list outside-acl permitted tcp outside/1.1.1.1(12345) ->
inside/192.168.1.1(1357) hit-cnt 1 (first hit)
```

Although 20 additional packets for this connection arrive on the outside interface, the traffic does not have to be checked against the ACL, and the hit count does not increase.

If one or more connections by the same host are initiated within the specified 10-minute interval (and the source and destination ports remain the same), then the hit count is incremented by 1, and the following syslog message displays at the end of the 10-minute interval:

```
%ASA|PIX-7-106100: access-list outside-acl permitted tcp outside/1.1.1.1(12345)->
inside/192.168.1.1(1357) hit-cnt 2 (600-second interval)
```

When the third ACE denies a packet, the ASA generates the following syslog message:

```
%ASA|PIX-2-106100: access-list outside-acl denied ip outside/3.3.3.3(12345) ->
inside/192.168.1.1(1357) hit-cnt 1 (first hit)
```

If 20 additional attempts occur within a 5-minute interval (the default), the following syslog message appears at the end of 5 minutes:

```
%ASA|PIX-2-106100: access-list outside-acl denied ip outside/3.3.3.3(12345) ->
inside/192.168.1.1(1357) hit-cnt 21 (300-second interval)
```

Feature History for ACL Logging

Table 23-2 lists the release history for this feature.

Table 23-2 Feature History for ACL Logging

Feature Name	Releases	Feature Information
ACL logging	7.0(1)	You can enable logging using syslog message 106100, which provides statistics for each ACE and lets you limit the number of syslog messages produced. We introduced the following command: access-list .
ACL Timestamp	8.3(1)	The ASA reports the timestamp for the last access rule hit.

Managing Deny Flows

This section includes the following topics:

- [Information About Managing Deny Flows, page 23-6](#)
- [Licensing Requirements for Managing Deny Flows, page 23-6](#)
- [Guidelines and Limitations, page 23-6](#)
- [Managing Deny Flows, page 23-7](#)
- [Monitoring Deny Flows, page 23-7](#)
- [Feature History for Managing Deny Flows, page 23-8](#)

Information About Managing Deny Flows

When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA has a maximum of 32 K logging flows for ACEs. A large number of flows can exist concurrently at any point of time. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent *deny* flows; the limit is placed on deny flows only (not on permit flows) because they can indicate an attack. When the limit is reached, the ASA does not create a new deny flow for logging until the existing flows expire.

For example, if someone initiates a DoS attack, the ASA can create a large number of deny flows in a short period of time. Restricting the number of deny flows prevents unlimited consumption of memory and CPU resources.

When you reach the maximum number of deny flows, the ASA issues syslog message 106100:

```
%ASA|PIX-1-106101: The number of ACL log deny-flows has reached limit (number).
```

The **access-list alert-interval** command sets the time interval for generating syslog message 106001. Syslog message 106001 alerts you that the ASA has reached a deny flow maximum. When the deny flow maximum is reached, another syslog message 106001 is generated if at least six seconds have passed since the last 106001 message was generated.

Licensing Requirements for Managing Deny Flows

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 and transparent firewall modes.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

The ASA places a limit on the number of concurrent *deny* flows only—not permit flows.

Default Settings

Table 23-1 lists the default settings for managing deny flows.

Table 23-3 Default Parameters for Managing Deny Flows

Parameters	Default
<i>numbers</i>	The <i>numbers</i> argument specifies the maximum number of deny flows. The default is 4096.
<i>secs</i>	The <i>secs</i> argument specifies the time, in seconds, between syslog messages. The default is 300.

Managing Deny Flows

To configure the maximum number of deny flows and to set the interval between deny flow alert messages (106100), enter the following command:

Command	Purpose
<code>access-list deny-flow-max <i>number</i></code>	Sets the maximum number of deny flows.
Example: <code>ciscoasa(config)# access-list deny-flow-max 3000</code>	The <i>numbers</i> argument specifies the maximum number, which can be between 1 and 4096. The default is 4096.

To set the amount of time between syslog messages (number 106101), which identifies that the maximum number of deny flows was reached, enter the following command:

Command	Purpose
<code>access-list alert-interval <i>secs</i></code>	Sets the time, in seconds, between syslog messages.
Example: <code>ciscoasa(config)# access-list alert-interval 200</code>	The <i>secs</i> argument specifies the time interval between each deny flow maximum message. Valid values are from 1 to 3600 seconds. The default is 300 seconds.

Monitoring Deny Flows

To monitor ACLs, enter one of the following commands:

Command	Purpose
<code>show access-list</code>	Displays access list entries by number.
<code>show running-config access-list</code>	Displays the current running ACL configuration.

Feature History for Managing Deny Flows

Table 23-2 lists the release history for this feature.

Table 23-4 Feature History for Managing Deny Flows

Feature Name	Releases	Feature Information
Managing Deny Flows	7.0(1)	You can configure the maximum number of deny flows and set the interval between deny flow alert messages. We introduced the following commands: access-list deny-flow and access-list alert-interval .



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 interoperability 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 28-3](#).

- 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:

```
ciscoasa# 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
```

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

* If the administrative distance is 255, the router does not believe the source of that route and does not install the route in the routing table.

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:

```
ciscoasa# 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 versus 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.

**Note**

Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

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.

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 8, “Configuring a Cluster of ASAs.”](#)

Dynamic Routing in Multiple Context Mode

In multiple context mode, each context maintains a separate routing table and routing protocol databases. This enables you to configure OSPFv2 and EIGRP independently in each context. You can configure EIGRP in some contexts and OSPFv2 in the same or different contexts. In mixed context mode, you can enable any of the dynamic routing protocols in contexts that are in routed mode. RIP and OSPFv3 are not supported in multiple context mode.

The following table lists the attributes for EIGRP, OSPFv2, route maps used for distributing routes into OSPFv2 and EIGRP processes, and prefix lists used in OSPFv2 to filter the routing updates entering or leaving an area when they are used in multiple context mode:

EIGRP	OSPFv2	Route Maps and Prefix Lists
One instance is supported per context.	Two instances are supported per context.	N/A
It is disabled in the system context.		N/A
Two contexts may use the same or different autonomous system numbers.	Two contexts may use the same or different area IDs.	N/A
Shared interfaces in two contexts may have multiple EIGRP instances running on them.	Shared interfaces in two contexts may have multiple OSPF instances running on them.	N/A
The interaction of EIGRP instances across shared interfaces is supported.	The interaction of OSPFv2 instances across shared interfaces is supported.	N/A

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 command:

Command	Purpose
<code>sysopt noproxyarp interface</code>	Disables proxy ARP requests.
<p>Example:</p> <pre>ciscoasa(config)# sysopt noproxyarp exampleinterface</pre>	



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 25-6](#)
- [Configuration Examples for Static or Default Routes, page 25-8](#)
- [Feature History for Static and Default Routes, page 25-9](#)

Information About Static and Default Routes

To route traffic to a nonconnected host or network, you must define a static route to the host or network or, at a minimum, a default route for any networks to which the ASA is not directly connected; for example, when there is a router between a network and the ASA.

Without a static or default route defined, traffic to nonconnected hosts or networks generates the following syslog message:

```
%ASA-6-110001: No route to dest_address from source_address
```

You might want to use static routes in single context mode in the following cases:

- Your networks use a different router discovery protocol from EIGRP, RIP, or OSPF.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.

The simplest option is to configure a default route to send all traffic to an upstream router, relying on the router to route the traffic for you. However, in some cases the default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the ASA.

In transparent firewall mode, for traffic that originates on the ASA and is destined for a nondirectly connected network, you need to configure either a default route or static routes so the ASA knows out of which interface to send traffic. Traffic that originates on the ASA might include communications to a

syslog server, Websense or N2H2 server, or AAA server. If you have servers that cannot all be reached through a single default route, then you must configure static routes. Additionally, the ASA supports up to three equal cost routes on the same interface for load balancing.

Licensing Requirements for Static and Default Routes

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6.

Failover Guidelines

Supports Stateful Failover of dynamic routing protocols.

Additional Guidelines

- IPv6 static routes are not supported in transparent mode in ASDM.
- In clustering, static route monitoring is only supported on the master unit. For information about clustering, see [Chapter 8, “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 25-4](#)
- [Configuring IPv6 Default and Static Routes, page 25-5](#)

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 section:

- [Adding or Editing a Static Route, page 25-3](#)

Adding or Editing a Static Route

To add or edit a static route, enter the following command:

Command	Purpose
<pre>route if_name dest_ip mask gateway_ip [distance]</pre> <p>Example: <pre>ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.1 [1]</pre></p>	<p>Enables you to add a static route.</p> <p>The <i>dest_ip</i> and <i>mask</i> arguments indicate the IP address for the destination network and the <i>gateway_ip</i> 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.</p> <p>The <i>distance</i> 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.</p> <p>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.</p>

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.

```
ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.1
ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.1.2
ciscoasa(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
<pre>route if_name 0.0.0.0 0.0.0.0 gateway_ip [distance tunneled]</pre> <p>Example:</p> <pre>ciscoasa(config)# route outside 0 0 192.168.2.4 tunneled</pre>	<p>Enables you to add a static route.</p> <p>The <i>dest_ip</i> and <i>mask</i> arguments indicate the IP address for the destination network and the <i>gateway_ip</i> 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.</p> <p>The <i>distance</i> 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.</p>

**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:

```
ciscoasa(config)# route outside 0 0 192.168.1 1
```

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	<pre>ipv6 route if_name ::/0 next_hop_ipv6_addr</pre> <p>Example:</p> <pre>ciscoasa(config)# ipv6 route inside 7fff::0/32 3FFE:1100:0:CC00::1</pre>	<p>Adds a default IPv6 route.</p> <p>The example routes packets for network 7fff::0/32 to a networking device on the inside interface at 3FFE:1100:0:CC00::1</p> <p>The address ::/0 is the IPv6 equivalent of any.</p>
Step 2	<pre>ipv6 route if_name destination next_hop_ipv6_addr [admin_distance]</pre> <p>Example:</p> <pre>ciscoasa(config)# ipv6 route inside 7fff::0/32 3FFE:1100:0:CC00::1 [110]</pre>	<p>Adds an IPv6 static route to the IPv6 routing table.</p> <p>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.</p>

**Note**

The **ipv6 route** command works the same way as the **route** command does, which is used to define IPv4 static routes.

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	<p>sla monitor <i>sla_id</i></p> <p>Example: <pre>ciscoasa(config)# sla monitor sla_id</pre></p>	<p>Configures the tracked object monitoring parameters by defining the monitoring process.</p> <p>If you are configuring a new monitoring process, you enter sla monitor configuration mode.</p> <p>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.</p>
Step 2	<p>type echo protocol ipIcmpEcho <i>target_ip</i> interface <i>if_name</i></p> <p>Example: <pre>ciscoasa(config-sla-monitor)# type echo protocol ipIcmpEcho target_ip interface if_name</pre></p>	<p>Specifies the monitoring protocol.</p> <p>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.</p> <p>The <i>target_ip</i> 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.</p>
Step 3	<p>sla monitor schedule <i>sla_id</i> [life {forever <i>seconds</i>}] [start-time {<i>hh:mm</i> [:<i>ss</i>] [<i>month day</i> <i>day month</i>] <i>pending</i> <i>now</i> after <i>hh:mm:ss</i>}] [ageout <i>seconds</i>] [recurring]</p> <p>Example: <pre>ciscoasa(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]</pre></p>	<p>Schedules the monitoring process.</p> <p>Typically, you will use the sla monitor schedule <i>sla_id</i> life forever start-time now command for the monitoring schedule, and allow the monitoring configuration to determine how often the testing occurs.</p> <p>However, you can schedule this monitoring process to begin in the future and to only occur at specified times.</p>
Step 4	<p>track <i>track_id</i> rtr <i>sla_id</i> reachability</p> <p>Example: <pre>ciscoasa(config)# track track_id rtr sla_id reachability</pre></p>	<p>Associates a tracked static route with the SLA monitoring process.</p> <p>The <i>track_id</i> argument is a tracking number you assign with this command. The <i>sla_id</i> argument is the ID number of the SLA process.</p>
Step 5	<p>Do one of the following to define the static route to be installed in the routing table while the tracked object is reachable.</p> <p>These options allow you to track a static route or a default route obtained through DHCP or PPPOE.</p> <p>route <i>if_name</i> <i>dest_ip</i> <i>mask</i> <i>gateway_ip</i> [<i>admin_distance</i>] track <i>track_id</i></p> <p>Example: <pre>ciscoasa(config)# route if_name dest_ip mask gateway_ip [admin_distance] track track_id</pre></p>	<p>Tracks a static route.</p> <p>You cannot use the tunneled option with the route command in static route tracking.</p>

Command	Purpose
Example: <pre>ciscoasa(config)# interface phy_if ciscoasa(config-if)# dhcp client route track track_id ciscoasa(config-if)# ip address dhcp setroute ciscoasa(config-if)# exit</pre>	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: <pre>ciscoasa(config)# interface phy_if ciscoasa(config-if)# pppoe client route track track_id ciscoasa(config-if)# ip address pppoe setroute ciscoasa(config-if)# exit</pre>	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.

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:

```
ciscoasa(config)# route inside 10.1.1.0 255.255.255.0 10.1.2.45 1
ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.1
ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.2
ciscoasa(config)# route outside 10.10.10.0 255.255.255.0 192.168.2.3
ciscoasa(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:

```
ciscoasa(config)# route inside 10.1.1.0 255.255.255.0 10.1.2.45 1
```

Feature History for Static and Default Routes

Table 25-1 lists each feature change and the platform release in which it was implemented.

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.
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 26-4](#)
- [Configuration Example for Route Maps, page 26-6](#)
- [Feature History for Route Maps, page 26-6](#)

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 recommend 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-igrp** 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 selects 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 allow 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
<pre>route-map name {permit deny} [sequence_number]</pre> <p>Example: ciscoasa(config)# route-map name {permit} [12]</p>	<p>Creates the route map entry. Enters route-map configuration mode.</p> <p>Route map entries are read in order. You can identify the order using the <i>sequence_number</i> argument, or the ASA uses the order in which you add route map entries.</p>

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 26-4](#)
- [Configuring the Metric Values for a Route Action, page 26-5](#)

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	<pre>route-map name {permit deny} [sequence_number]</pre> <p>Example: ciscoasa(config)# route-map name {permit} [12]</p>	<p>Creates the route map entry. Enters route-map configuration mode.</p> <p>Route map entries are read in order. You can identify the order using the <i>sequence_number</i> option, or the ASA uses the order in which you add route map entries.</p>
Step 2	<p>Enter one of the following match commands to match routes to a specified destination address:</p> <pre>match ip address acl_id [acl_id] [...] [prefix-list]</pre> <p>Example: ciscoasa(config-route-map)# match ip address acl_id [acl_id] [...]</p>	<p>Matches any routes that have a destination network that matches a standard ACL or prefix list.</p> <p>If you specify more than one ACL, then the route can match any of the ACLs.</p>

Command	Purpose
match metric <i>metric_value</i> Example: ciscoasa(config-route-map)# match metric 200	Matches any routes that have a specified metric. The <i>metric_value</i> can range from 0 to 4294967295.
match ip next-hop <i>acl_id</i> [<i>acl_id</i>] [...]	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 <i>if_name</i> Example: ciscoasa(config-route-map)# match interface <i>if_name</i>	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 <i>acl_id</i> [<i>acl_id</i>] [...]	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]}	Matches the route type.
Example: ciscoasa(config-route-map)# match route-type internal type-1	

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 <i>name</i> { permit deny } [<i>sequence_number</i>]	Creates the route map entry. Enters route-map configuration mode. Route map entries are read in order. You can identify the order using the <i>sequence_number</i> argument, or the ASA uses the order in which you add route map entries.
	Example: ciscoasa(config)# route-map <i>name</i> {permit} [12]	
Step 2	To set a metric for the route map, enter one or more of the following set commands:	

Command	Purpose
set metric <i>metric_value</i> Example: ciscoasa(config-route-map)# set metric 200	Sets the metric value. The <i>metric_value</i> argument can range from 0 to 294967295.
set metric-type { type-1 type-2 } Example: ciscoasa(config-route-map)# set metric-type type-2	Sets the metric type. The <i>metric-type</i> argument can be type-1 or type-2.

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.

```
ciscoasa(config)# route-map 1-to-2 permit
ciscoasa(config-route-map)# match metric 1
ciscoasa(config-route-map)# set metric 5
ciscoasa(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:

```
ciscoasa(config)# route outside 10.1.1.0 255.255.255.0 192.168.1.1
ciscoasa(config-route-map)# access-list mymap2 line 1 permit 10.1.1.0 255.255.255.0
ciscoasa(config-route-map)# route-map mymap2 permit 10
ciscoasa(config-route-map)# match ip address mymap2
ciscoasa(config-route-map)# router eigrp 1
ciscoasa(config)# redistribute static metric 250 250 1 1 1 route-map mymap2
```

Feature History for Route Maps

Table 26-1 lists each feature change and the platform release in which it was implemented.

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 .
Enhanced support for static and dynamic route maps	8.0(2)	Enhanced support for dynamic and static route maps was added.

Table 26-1 Feature History for Route Maps (continued)

Feature Name	Platform Releases	Feature Information
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 8, “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 <i>ospf process_id</i> Example: ciscoasa(config)# router ospf 2	Creates an OSPF routing process and enters router configuration mode for this OSPF process. The <i>process_id</i> 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 <i>ip_address mask area area_id</i> Example: ciscoasa(config)# router ospf 2 ciscoasa(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	<pre>router ospf process_id</pre> <p>Example: ciscoasa(config)# router ospf 2</p>	<p>Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute.</p> <p>The <i>process_id</i> 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.</p>
Step 2	Do one of the following to redistribute the selected route type into the OSPF routing process:	

Command	Purpose
<pre> redistribute connected [[metric <i>metric-value</i>] [metric-type {type-1 type-2}] [tag <i>tag_value</i>] [subnets] [route-map <i>map_name</i>] </pre> <p>Example:</p> <pre> ciscoasa(config)# redistribute connected 5 type-1 route-map-practice </pre>	<p>Redistributes connected routes into the OSPF routing process.</p>
<pre> redistribute static [metric <i>metric-value</i>] [metric-type {type-1 type-2}] [tag <i>tag_value</i>] [subnets] [route-map <i>map_name</i>] </pre> <p>Example:</p> <pre> ciscoasa(config)# redistribute static 5 type-1 route-map-practice </pre>	<p>Redistributes static routes into the OSPF routing process.</p>
<pre> redistribute ospf pid [match {internal external [1 2] nssa-external [1 2]}] [metric <i>metric-value</i>] [metric-type {type-1 type-2}] [tag <i>tag_value</i>] [subnets] [route-map <i>map_name</i>] </pre> <p>Example:</p> <pre> ciscoasa(config)# route-map 1-to-2 permit ciscoasa(config-route-map)# match metric 1 ciscoasa(config-route-map)# set metric 5 ciscoasa(config-route-map)# set metric-type type-1 ciscoasa(config-route-map)# router ospf 2 ciscoasa(config-rtr)# redistribute ospf 1 route-map 1-to-2 </pre>	<p>Allows you to redistribute routes from an OSPF routing process into another OSPF routing process.</p> <p>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.</p> <p>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.</p>

Command	Purpose
<pre>redistribute rip [metric metric-value] [metric-type {type-1 type-2}] [tag tag_value] [subnets] [route-map map_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# redistribute rip 5 ciscoasa(config-route-map)# match metric 1 ciscoasa(config-route-map)# set metric 5 ciscoasa(config-route-map)# set metric-type type-1 ciscoasa(config-rtr)# redistribute ospf 1 route-map 1-to-2</pre>	Allows you to redistribute routes from a RIP routing process into the OSPF routing process.
<pre>redistribute eigrp as-num [metric metric-value] [metric-type {type-1 type-2}] [tag tag_value] [subnets] [route-map map_name]</pre> <p>Example:</p> <pre>ciscoasa(config)# redistribute eigrp 2 ciscoasa(config-route-map)# match metric 1 ciscoasa(config-route-map)# set metric 5 ciscoasa(config-route-map)# set metric-type type-1 ciscoasa(config-rtr)# redistribute ospf 1 route-map 1-to-2</pre>	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	<code>router ospf process_id</code> Example: <code>ciscoasa(config)# router ospf 1</code>	Creates an OSPF routing process and enters router configuration mode for this OSPF process. The <i>process_id</i> 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	<code>summary-address ip_address mask</code> <code>[not-advertise] [tag tag]</code> Example: <code>ciscoasa(config)# router ospf 1</code> <code>ciscoasa(config-rtr)# summary-address</code> <code>10.1.0.0 255.255.0.0</code>	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	<pre>router ospf process_id</pre> <p>Example:</p> <pre>ciscoasa(config)# router ospf 1</pre>	<p>Creates an OSPF routing process and enters router configuration mode for this OSPF process.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<pre>area area-id range ip-address mask [advertise not-advertise]</pre> <p>Example:</p> <pre>ciscoasa(config)# router ospf 1 ciscoasa(config-rtr)# area 17 range 12.1.0.0 255.255.0.0</pre>	<p>Sets the address range.</p> <p>In this example, the address range is set between OSPF areas.</p>

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 <i>ospf process_id</i> Example: ciscoasa(config)# router ospf 2	Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute. The <i>process_id</i> 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 <i>ip_address mask area area_id</i> Example: ciscoasa(config)# router ospf 2 ciscoasa(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 <i>interface_name</i> Example: ciscoasa(config)# interface <i>my_interface</i>	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: ciscoasa(config-interface)# ospf authentication message-digest	Specifies the authentication type for an interface.

Command	Purpose
<p>ospf authentication-key <i>key</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf authentication-key cisco</pre></p>	<p>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.</p> <p>The <i>key</i> argument can be any continuous string of characters up to 8 bytes in length.</p> <p>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.</p>
<p>ospf cost <i>cost</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf cost 20</pre></p>	<p>Allows you to explicitly specify the cost of sending a packet on an OSPF interface. The <i>cost</i> is an integer from 1 to 65535.</p> <p>In this example, the cost is set to 20.</p>
<p>ospf dead-interval <i>seconds</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf dead-interval 40</pre></p>	<p>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.</p> <p>In this example, the dead interval is set to 40.</p>
<p>ospf hello-interval <i>seconds</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf hello-interval 10</pre></p>	<p>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.</p> <p>In this example, the hello interval is set to 10.</p>
<p>ospf message-digest-key <i>key_id md5 key</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf message-digest-key 1 md5 cisco</pre></p>	<p>Enables OSPF MD5 authentication.</p> <p>The following argument values can be set:</p> <ul style="list-style-type: none"> <i>key_id</i>—An identifier in the range from 1 to 255. <i>key</i>—An alphanumeric password of up to 16 bytes. <p>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.</p> <p>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.</p>
<p>ospf priority <i>number_value</i></p> <p>Example: <pre>ciscoasa(config-interface)# ospf priority 20</pre></p>	<p>Allows you to set the priority to help determine the OSPF designated router for a network.</p> <p>The <i>number_value</i> argument ranges from 0 to 255.</p> <p>In this example, the priority number value is set to 20.</p>

Command	Purpose
ospf retransmit-interval <i>seconds</i> Example: <pre>ciscoasa(config-interface)# ospf retransmit-interval seconds</pre>	<p>Allows you to specify the number of seconds between LSA retransmissions for adjacencies belonging to an OSPF interface.</p> <p>The value for <i>seconds</i> 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.</p> <p>In this example, the retransmit-interval value is set to 15.</p>
ospf transmit-delay <i>seconds</i> Example: <pre>ciscoasa(config-interface)# ospf transmit-delay 5</pre>	<p>Sets the estimated number of seconds required to send a link-state update packet on an OSPF interface. The <i>seconds</i> value ranges from 1 to 65535 seconds. The default value is 1 second.</p> <p>In this example, the transmit-delay is 5 seconds.</p>
ospf network point-to-point non-broadcast Example: <pre>ciscoasa(config-interface)# ospf network point-to-point non-broadcast</pre>	<p>Specifies the interface as a point-to-point, non-broadcast network.</p> <p>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.</p>

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 <i>process_id</i> Example: ciscoasa(config)# router ospf 2	Creates an OSPF routing process and enters router configuration mode for the OSPF process that you want to redistribute. The <i>process_id</i> 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 <i>area-id</i> authentication Example: ciscoasa(config-rtr)# area 0 authentication	Enables authentication for an OSPF area.
	area <i>area-id</i> authentication message-digest Example: ciscoasa(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 clear 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	<pre>router ospf process_id</pre> <p>Example: ciscoasa(config)# router ospf 2</p>	<p>Creates an OSPF routing process and enters router configuration mode for the OSPF routing process that you want to redistribute.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<p>Do one of the following to configure optional OSPF NSSA parameters:</p> <pre>area area-id nssa [no-redistribution] [default-information-originate]</pre> <p>Example: ciscoasa(config-rtr)# area 0 nssa</p> <pre>summary-address ip_address mask [not-advertise] [tag tag]</pre> <p>Example: ciscoasa(config-rtr)# summary-address 10.1.0.0 255.255.0.0</p>	<p>Defines an NSSA area.</p> <p>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.</p> <p>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.</p>



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
<pre>router-id cluster-pool hostname A.B.C.D ip_pool</pre> <p>Example:</p> <pre>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</pre>	<p>Specifies the router ID cluster pool for Layer 3 clustering.</p> <p>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 <i>ip_pool</i> argument specifies the name of the IP address pool.</p> <p>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.</p>

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	<code>router ospf process_id</code> Example: <code>ciscoasa(config)# router ospf 2</code>	Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process. The <i>process_id</i> 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	<code>neighbor addr [interface if_name]</code> Example: <code>ciscoasa(config-rtr)# neighbor 255.255.0.0 [interface my_interface]</code>	Defines the OSPFv2 neighborhood. The <i>addr</i> argument is the IP address of the OSPFv2 neighbor. The <i>if_name</i> 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	<code>router ospf process_id</code> Example: <code>ciscoasa(config)# router ospf 2</code>	Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process. The <i>process_id</i> 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	<code>timers spf spf-delay spf-holdtime</code> Example: <code>ciscoasa(config-rtr)# timers spf 10 120</code>	Configures the route calculation times. The <i>spf-delay</i> 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. The <i>spf-holdtime</i> 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.

- 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	<pre>router ospf process_id</pre> <p>Example: ciscoasa(config)# router ospf 2 </p>	<p>Creates an OSPFv2 routing process and enters router configuration mode for this OSPFv2 process.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<pre>log-adj-changes [detail]</pre> <p>Example: ciscoasa(config-rtr)# log-adj-changes [detail] </p>	<p>Configures logging for neighbors going up or down.</p>

- 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.
- 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
Command <code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config)# ipv6 router ospf 10</code>	Creates an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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	<code>interface interface_name</code> Example: <code>ciscoasa(config)# interface GigabitEthernet0/0</code>	Enables an interface.
Step 2	<code>ipv6 ospf process-id area area_id</code> Example: <code>ciscoasa(config)# ipv6 ospf 200 area 100</code>	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	<pre>ipv6 router ospf <i>process-id</i></pre> <p>Example: <pre>ciscoasa(config-if)# ipv6 router ospf 10</pre></p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	<pre>ipv6 ospf area [<i>area-num</i>] [<i>instance</i>]</pre> <p>Example: <pre>ciscoasa(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</pre></p>	<p>Creates an OSPFv3 area.</p> <p>The <i>area-num</i> 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.</p>
Step 3	<p>Do one of the following to configure OSPFv3 interface parameters:</p> <pre>ipv6 ospf cost <i>interface-cost</i></pre> <p>Example: <pre>ciscoasa(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</pre></p>	<p>Explicitly specifies the cost of sending a packet on an interface. The <i>interface-cost</i> 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.</p>

Command	Purpose
<p>ipv6 ospf database-filter all out</p> <p>Example:</p> <pre>ciscoasa(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</pre>	<p>Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.</p>
<p>ipv6 ospf dead-interval seconds</p> <p>Example:</p> <pre>ciscoasa(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</pre>	<p>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.</p>

Command	Purpose
<pre> ipv6 ospf encryption {ipsec spi spi esp encryption-algorithm [[key-encryption-type] key] authentication-algorithm [[key-encryption-type] key null} Example: ciscoasa(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 </pre>	<p>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 <i>encryption-algorithm</i> argument specifies the encryption algorithm to be used with ESP. Valid values include the following:</p> <ul style="list-style-type: none"> • aes-cdc—Enables AES-CDC encryption. • 3des—Enables 3DES encryption. • des—Enables DES encryption. • null—Specifies ESP with no encryption. <p>The <i>key-encryption-type</i> argument can be one of the following two values:</p> <ul style="list-style-type: none"> • 0—The key is not encrypted. • 7—The key is encrypted. <p>The <i>key</i> 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 <i>authentication-algorithm</i> argument specifies the encryption authentication algorithm to be used, which can be one of the following:</p> <ul style="list-style-type: none"> • md5—Enables message digest 5 (MD5). • sha1—Enables SHA-1. <p>The null keyword overrides area encryption.</p> <p>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.</p>
<pre> ipv6 ospf flood-reduction Example: ciscoasa(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 </pre>	<p>Specifies the flood reduction of LSAs to the interface.</p>

Command	Purpose
<pre> ipv6 ospf hello-interval <i>seconds</i> Example: ciscoasa(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 </pre>	<p>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.</p>
<pre> ipv6 ospf mtu-ignore Example: ciscoasa(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 </pre>	<p>Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.</p>
<pre> ipv6 ospf network {broadcast point-to-point non-broadcast} Example: ciscoasa(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 </pre>	<p>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.</p>

Command	Purpose
<pre> ipv6 ospf priority <i>number-value</i> Example: ciscoasa(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 </pre>	<p>Sets the router priority, which helps determine the designated router for a network. Valid values range from 0 to 255.</p>
<pre> ipv6 ospf neighbor <i>ipv6-address</i> [priority <i>number</i>] [poll-interval <i>seconds</i>] [cost <i>number</i>] [database-filter all out] Example: ciscoasa(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 </pre>	<p>Configures OSPFv3 router interconnections to non-broadcast networks.</p>

Command	Purpose
<p>Command</p> <pre>ipv6 ospf retransmit-interval seconds</pre> <p>Example:</p> <pre>ciscoasa(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</pre>	<p>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.</p>
<p>Command</p> <pre>ipv6 ospf transmit-delay seconds</pre> <p>Example:</p> <pre>ciscoasa(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</pre>	<p>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.</p>

- 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 <code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config)# ipv6 router ospf 10</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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:	
<code>area</code> Example: <code>ciscoasa(config-rtr)# area 10</code>	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: <code>ciscoasa(config-rtr)# default originate</code>	Sets a command to its default value. The originate parameter distributes the default route.
default-information Example: <code>ciscoasa(config-rtr)# default-information</code>	Controls distribution of default information.
distance Example: <code>ciscoasa(config-rtr)# distance 200</code>	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.
<code>exit</code> Example: <code>ciscoasa(config-rtr)# exit</code>	Exits from IPv6 router configuration mode.
ignore Example: <code>ciscoasa(config-rtr)# ignore lsa</code>	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: <code>ciscoasa(config-rtr)# log-adjacency-changes detail</code>	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
<p>passive-interface <i>[interface_name]</i></p> <p>Example: <pre>ciscoasa(config-rtr)# passive-interface inside</pre></p>	<p>Suppresses the sending and receiving of routing updates on an interface. The <i>interface_name</i> argument specifies the name of the interface on which the OSPFv3 process is running.</p>
<p>redistribute</p> <p>Example: <pre>ciscoasa(config-rtr)# redistribute ospf</pre></p>	<p>Configures the redistribution of routes from one routing domain into another according to the following parameters:</p> <ul style="list-style-type: none"> • connected—Specifies connected routes. • ospf—Specifies OSPFv3 routes. • static—Specifies static routes.
<p>router-id</p> <p>Example: <pre>ciscoasa(config-rtr)# router-id 10.1.1.1</pre></p>	<p>Creates a fixed router ID for a specified process with the following parameters:</p> <ul style="list-style-type: none"> • 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.
<p>summary-prefix</p> <p>Example: <pre>ciscoasa(config-if)# ipv6 router ospf 1 ciscoasa(config-router)# router-id 192.168.3.3 ciscoasa(config-router)# summary-prefix FECO::/24 ciscoasa(config-router)# redistribute static</pre></p>	<p>Configures IPv6 address summaries with valid values from 0 to 128. The X:X:X:X::X/ parameter specifies the IPv6 prefix.</p>
<p>timers</p> <p>Example: <pre>ciscoasa(config)# ipv6 router ospf 10 ciscoasa(config-rtr)# timers throttle spf 6000 12000 14000</pre></p>	<p>Adjusts routing timers. The routing timer parameters are the following:</p> <ul style="list-style-type: none"> • 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	Do one of the following to configure optional OSPFv3 area parameters:	
	<pre>area area-id default-cost cost</pre> <p>Example: ciscoasa(config-rtr)# area 1 default-cost nssa</p>	<p>Sets the summary default cost of an NSSA area or a stub area.</p>
	<pre>area area-id range ipv6-prefix/ prefix-length [advertise not advertise] [cost cost]</pre> <p>Example: ciscoasa(config-rtr)# area 1 range FE01:1::1/64</p>	<p>Summarizes routes that match the address and mask for border routers only.</p> <p>The <i>area-id</i> argument identifies the area for which routes are to be summarized. The value can be specified as a decimal or an IPv6 prefix. The <i>ipv6-prefix</i> argument specifies the IPv6 prefix. The <i>prefix-length</i> 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.</p>
	<pre>area area-id nssa</pre> <p>Example: ciscoasa(config-rtr)# area 1 nssa</p>	<p>Specifies an NSSA area.</p>

Command	Purpose
<p>area <i>area-id</i> stub</p> <p>Example: ciscoasa(config-rtr)# area 1 stub</p>	Specifies a stub area.
<p>area <i>area-id</i> virtual-link <i>router-id</i> [hello-interval <i>seconds</i>] [retransmit-interval <i>seconds</i>] [transmit-delay <i>seconds</i>] [dead-interval <i>seconds</i>] [ttl-security hops <i>hop-count</i>]</p> <p>Example: ciscoasa(config-rtr)# area 1 virtual-link 192.168.255.1 hello-interval 5</p>	<p>Defines a virtual link and its parameters.</p> <p>The <i>area-id</i> argument identifies the area for which routes are to be summarized. The virtual link keyword specifies the creation of a virtual link neighbor. The <i>router-id</i> argument specifies the router ID that is associated with the virtual link neighbor. Enter the show ospf or show ipv6 ospf 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 <i>seconds</i> 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 <i>seconds</i> 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 <i>seconds</i> 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 <i>hop-count</i> argument value can range from 1 to 254.</p>

- 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	<code>ipv6 router ospf <i>process_id</i></code> Example: <code>ciscoasa(config-if)# ipv6 router ospf 1</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process_id</i> 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	<code>passive-interface [<i>interface_name</i>]</code> Example: <code>ciscoasa(config-rtr)# passive-interface inside</code>	Suppresses the sending and receiving of routing updates on an interface. The <i>interface_name</i> argument specifies the name of the interface on which the OSPFv3 process is running. If the <i>no interface_name</i> argument is specified, all of the interfaces in the OSPFv3 process <i>process_id</i> 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	<pre>ipv6 router ospf process_id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<pre>distance [ospf {external inter-area intra-area}] distance</pre> <p>Example: ciscoasa(config-rtr)# distance ospf external 200</p>	<p>Sets the administrative distance for OSPFv3 routes.</p> <p>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 OSPFv3. The intra-area keyword specifies the intra-area routes for OSPFv3. The <i>distance</i> argument specifies the administrative distance, which is an integer from 10 to 254.</p>

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 OSPF 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 IPv6 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	<code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config-if)# ipv6 router ospf 1</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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	<code>timers lsa arrival milliseconds</code> Example: <code>ciscoasa(config-rtr)# timers lsa arrival 2000</code>	Sets the minimum interval at which the ASA accepts the same LSA from OSPF neighbors. The <i>milliseconds</i> 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	<code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config-if)# ipv6 router ospf 1</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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	<code>timers pacing flood milliseconds</code> Example: <code>ciscoasa(config-rtr)# timers lsa flood 20</code>	Configures LSA flood packet pacing. The <i>milliseconds</i> 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	<code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config-if)# ipv6 router ospf 1</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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	<code>timers pacing lsa-group seconds</code> Example: <code>ciscoasa(config-rtr)# timers pacing lsa-group 300</code>	Changes the interval at which OSPFv3 LSAs are collected into a group and refreshed, checksummed, or aged. The <i>seconds</i> 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	<code>ipv6 router ospf process-id</code> Example: <code>ciscoasa(config-if)# ipv6 router ospf 1</code>	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process-id</i> 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	<code>timers pacing retransmission milliseconds</code> Example: <code>ciscoasa(config-rtr)# timers pacing retransmission 100</code>	Configures LSA retransmission packet pacing. The <i>milliseconds</i> 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	Choose one of the following options:	

Command	Purpose
<pre>timers throttle lsa milliseconds1 milliseconds2 milliseconds3</pre> <p>Example:</p> <pre>ciscoasa(config-rtr)# timers throttle lsa 500 6000 8000</pre>	<p>Configures OSPFv3 LSA throttling.</p> <p>The <i>milliseconds1</i> argument specifies the delay in milliseconds to generate the first occurrence of the LSA. The <i>milliseconds2</i> argument specifies the maximum delay in milliseconds to originate the same LSA. The <i>milliseconds3</i> argument specifies the minimum delay in milliseconds to originate the same LSA.</p> <p>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.</p> <p>The default values for LSA throttling are the following:</p> <ul style="list-style-type: none"> For <i>milliseconds1</i>, the default value is 0 milliseconds. For <i>milliseconds2</i> and <i>milliseconds3</i>, the default value is 5000 milliseconds.
<pre>timers throttle spf milliseconds1 milliseconds2 milliseconds3</pre> <p>Example:</p> <pre>ciscoasa(config-rtr)# timers throttle spf 5000 12000 16000</pre>	<p>Configures OSPFv3 SPF throttling.</p> <p>The <i>milliseconds1</i> argument specifies the delay in milliseconds to receive a change to the SPF calculation. The <i>milliseconds2</i> argument specifies the delay in milliseconds between the first and second SPF calculations. The <i>milliseconds3</i> argument specifies the maximum wait time in milliseconds for SPF calculations.</p> <p>For SPF throttling, if <i>milliseconds2</i> or <i>milliseconds3</i> is less than <i>milliseconds1</i>, then OSPFv3 automatically corrects to the <i>milliseconds1</i> value. Similarly, if <i>milliseconds3</i> is less than <i>milliseconds2</i>, then OSPFv3 automatically corrects to the <i>milliseconds2</i> value.</p> <p>The default values for SPF throttling are the following:</p> <ul style="list-style-type: none"> For <i>milliseconds1</i>, the default value is 5000 milliseconds. For <i>milliseconds2</i> and <i>milliseconds3</i>, 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	<pre>ipv6 ospf neighbor ipv6-address [priority number] [poll-interval seconds] [cost number] [database-filter all out]</pre> <p>Example: ciscoasa(config-if)# interface ethernet0/0 ipv6 ospf neighbor FE80::A8BB:CCFF:FE00:C01</p>	<p>Configures OSPFv3 router interconnections to non-broadcast networks.</p>

-
- 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<pre>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]</pre> <p>Example: ciscoasa(config-rtr)# default metric 5</p>	<p>Returns an optional parameter to its default value.</p> <p>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.</p>

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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	<pre>log-adjacency-changes [detail]</pre> <p>Example: ciscoasa(config-rtr)# log-adjacency-changes detail</p>	<p>Configures the router to send a syslog message when an OSPFv3 neighbor goes up or down.</p> <p>The detail keyword sends a syslog message for each state, not only when an OSPFv3 neighbor goes up or down.</p>

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 send 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	<pre>router ospf process_id</pre> <p>Example: ciscoasa(config-if)# router ospf 1</p>	<p>Enables an OSPFv2 routing process and enters router configuration mode.</p> <p>The <i>process_id</i> 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.</p>
Step 2	<pre>ignore lsa mospf</pre> <p>Example: ciscoasa(config-rtr)# ignore lsa mospf</p>	<p>Suppresses the sending of syslog messages when the router receives unsupported LSA Type 6 MOSPF packets.</p>

-
- 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
<pre>compatible rfc1583</pre> <p>Example: ciscoasa (config-rtr)# compatible rfc1583</p>	<p>Restores the methods that are used to calculate summary route costs according to RFC 1583.</p>

-
- 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	<pre>default-information originate [always] metric metric-value [metric-type type-value] [route-map map-name]</pre> <p>Example: ciscoasa(config-rtr)# default-information originate always metric 3 metric-type 2</p>	<p>Generates a default external route into an OSPFv3 routing domain.</p> <p>The always keyword advertises the default route whether or not the default route exists. The metric <i>metric-value</i> 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 <i>type-value</i> 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:</p> <ul style="list-style-type: none"> • 1—Type 1 external route • 2—Type 2 external route <p>The default is the type 2 external route. The route-map <i>map-name</i> keyword-argument pair specifies the routing process that generates the default route if the route map is satisfied.</p>

- 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:
- 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	<code>ipv6 router ospf process-id</code> Example: ciscoasa(config-if)# ipv6 router ospf 1	Enables an OSPFv3 routing process and enters IPv6 router configuration mode. The <i>process_id</i> 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	<code>summary-prefix prefix [not-advertise tag tag-value]</code> Example: ciscoasa(config-if)# ipv6 router ospf 1 ciscoasa(config-rtr)# router-id 192.168.3.3 ciscoasa(config-rtr)# summary-prefix FECO::/24 ciscoasa(config-rtr)# redistribute static	Configures an IPv6 summary prefix. The <i>prefix</i> 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	<pre>ipv6 router ospf process-id</pre> <p>Example: ciscoasa(config-if)# ipv6 router ospf 1</p>	<p>Enables an OSPFv3 routing process and enters IPv6 router configuration mode.</p> <p>The <i>process-id</i> 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.</p>
Step 2	<pre>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]</pre> <p>Example: ciscoasa(config-rtr)# redistribute connected 5 type-1</p>	<p>Redistributes IPv6 routes from one OSPFv3 process into another.</p> <p>The <i>source-protocol</i> argument specifies the source protocol from which routes are being redistributed, which can be static, connected, or OSPFv3. The <i>process-id</i> 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 <i>metric-value</i> 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 <i>type-value</i> 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 <i>map-tag</i> argument identifies a configured route map.</p>

Removing the OSPF Configuration

To remove the entire OSPFv2 configuration that you have already enabled, enter the following commandperform the following steps:

Command	Purpose
<pre>clear configure router ospf pid</pre> <p>Example: <pre>ciscoasa(config)# clear configure router ospf 1000</pre></p>	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
<pre>clear configure ipv6 router ospf process-id</pre> <p>Example: <pre>ciscoasa(config)# clear configure ipv6 router ospf 1000</pre></p>	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:


```
ciscoasa(config)# router ospf 2
ciscoasa(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:

```

ciscoasa(config)# route-map 1-to-2 permit
ciscoasa(config-route-map)# match metric 1
ciscoasa(config-route-map)# set metric 5
ciscoasa(config-route-map)# set metric-type type-1
ciscoasa(config-route-map)# router ospf 2
ciscoasa(config-rtr)# redistribute ospf 1 route-map 1-to-2

```

Step 3 (Optional) To configure OSPFv2 interface parameters, enter the following commands:

```

ciscoasa(config)# router ospf 2
ciscoasa(config-rtr)# network 10.0.0.0 255.0.0.0 area 0
ciscoasa(config-rtr)# interface inside
ciscoasa(config-interface)# ospf cost 20
ciscoasa(config-interface)# ospf retransmit-interval 15
ciscoasa(config-interface)# ospf transmit-delay 10
ciscoasa(config-interface)# ospf priority 20
ciscoasa(config-interface)# ospf hello-interval 10
ciscoasa(config-interface)# ospf dead-interval 40
ciscoasa(config-interface)# ospf authentication-key cisco
ciscoasa(config-interface)# ospf message-digest-key 1 md5 cisco
ciscoasa(config-interface)# ospf authentication message-digest

```

Step 4 (Optional) To configure OSPFv2 area parameters, enter the following commands:

```

ciscoasa(config)# router ospf 2
ciscoasa(config-rtr)# area 0 authentication
ciscoasa(config-rtr)# area 0 authentication message-digest
ciscoasa(config-rtr)# area 17 stub
ciscoasa(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:

```

ciscoasa(config-rtr)# timers spf 10 120
ciscoasa(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:

```

ciscoasa(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:

```
ciscoasa(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:
- Check the **Enable OSPFv3 Process** check box.
 - 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 **1122334455aabbccdde** 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:

```
ciscoasa (config)# interface GigabitEthernet3/1
ciscoasa (config-if)# ipv6 enable
ciscoasa (config-if)# ipv6 ospf 1 area 1
```

The following is sample output from the **show running-config ipv6** command:

```
ciscoasa (config)# show running-config ipv6
ipv6 router ospf 1
  log-adjacency-changes
```

The following is sample output from the **show running-config interface** command:


```
ciscoasa (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:

```
ciscoasa (config)# interface GigabitEthernet3/1
ciscoasa (config-if)# nameif fda
ciscoasa (config-if)# security-level 100
ciscoasa (config-if)# ip address 10.1.11.1 255.255.255.0 standby 10.1.11.2
ciscoasa (config-if)# ipv6 address 9098::10/64 standby 9098::11
ciscoasa (config-if)# ipv6 enable
ciscoasa (config-if)# ipv6 ospf cost 900
ciscoasa (config-if)# ipv6 ospf hello-interval 20
ciscoasa (config-if)# ipv6 ospf network broadcast
ciscoasa (config-if)# ipv6 ospf database-filter all out
ciscoasa (config-if)# ipv6 ospf flood-reduction
ciscoasa (config-if)# ipv6 ospf mtu-ignore
ciscoasa (config-if)# ipv6 ospf 1 area 1 instance 100
ciscoasa (config-if)# ipv6 ospf encryption ipsec spi 890 esp null md5
12345678901234567890123456789012

ciscoasa (config)# ipv6 router ospf 1
ciscoasa (config)# area 1 nssa
ciscoasa (config)# distance ospf intra-area 190 inter-area 100 external 100
ciscoasa (config)# timers lsa arrival 900
ciscoasa (config)# timers pacing flood 100
ciscoasa (config)# timers throttle lsa 900 900 900
ciscoasa (config)# passive-interface fda
ciscoasa (config)# log-adjacency-changes
ciscoasa (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
<code>show ospf [process-id [area-id]]</code>	Displays general information about OSPFv2 routing processes.
<code>show ospf border-routers</code>	Displays the internal OSPFv2 routing table entries to the ABR and ASBR.

Command	Purpose
<code>show ospf [process-id [area-id]] database</code>	Displays lists of information related to the OSPFv2 database for a specific router.
<code>show ospf flood-list if-name</code>	<p>Displays a list of LSAs waiting to be flooded over an interface (to observe OSPF v2packet pacing).</p> <p>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:</p> <ul style="list-style-type: none"> • 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. <p>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.</p>
<code>show ospf interface [if_name]</code>	Displays OSPFv2-related interface information.
<code>show ospf neighbor [interface-name] [neighbor-id] [detail]</code>	Displays OSPFv2 neighbor information on a per-interface basis.
<code>show ospf request-list neighbor if_name</code>	Displays a list of all LSAs requested by a router.
<code>show ospf retransmission-list neighbor if_name</code>	Displays a list of all LSAs waiting to be resent.
<code>show ospf [process-id] summary-address</code>	Displays a list of all summary address redistribution information configured under an OSPFv2 process.
<code>show ospf [process-id] traffic</code>	Displays a list of different types of packets being sent or received by a specific OSPFv2 instance.
<code>show ospf [process-id] virtual-links</code>	Displays OSPFv2-related virtual links information.
<code>show route cluster</code>	Displays additional OSPFv2 route synchronization information in clustering.

To monitor or display various OSPFv3 routing statistics, enter one of the following commands:

Command	Purpose
<code>show ipv6 ospf [process-id [area-id]]</code>	Displays general information about OSPFv3 routing processes.
<code>show ipv6 ospf [process-id] border-routers</code>	Displays the internal OSPFv3 routing table entries to the ABR and ASBR.
<code>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]</code>	Displays lists of information related to the OSPFv3 database for a specific router.
<code>show ipv6 ospf [process-id [area-id]] events</code>	Displays OSPFv3 event information.
<code>show ipv6 ospf [process-id] [area-id] flood-list interface-type interface-number</code>	<p>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:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>There are no configuration tasks for this feature; it occurs automatically.</p>
<code>show ipv6 ospf [process-id] [area-id] interface [type number] [brief]</code>	Displays OSPFv3-related interface information.
<code>show ipv6 ospf neighbor [process-id] [area-id] [interface-type interface-number] [neighbor-id] [detail]</code>	Displays OSPFv3 neighbor information on a per-interface basis.
<code>show ipv6 ospf [process-id] [area-id] request-list [neighbor] [interface] [interface-neighbor]</code>	Displays a list of all LSAs requested by a router.

Command	Purpose
<code>show ipv6 ospf [process-id] [area-id] retransmission-list [neighbor] [interface] [interface-neighbor]</code>	Displays a list of all LSAs waiting to be resent.
<code>show ipv6 ospf statistic [process-id] [detail]</code>	Displays various OSPFv3 statistics.
<code>show ipv6 ospf [process-id] summary-prefix</code>	Displays a list of all summary address redistribution information configured under an OSPFv3 process.
<code>show ipv6 ospf [process-id] timers [lsa-group rate-limit]</code>	Displays OSPFv3 timers information.
<code>show ipv6 ospf [process-id] traffic [interface_name]</code>	Displays OSPFv3 traffic-related statistics.
<code>show ipv6 ospf virtual-links</code>	Displays OSPFv3-related virtual links information.
<code>show ipv6 route cluster [failover] [cluster] [interface] [ospf] [summary]</code>	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](#)
- [Configuring EIGRP, page 28-3](#)
- [Customizing EIGRP, page 28-5](#)
- [Monitoring EIGRP, page 28-18](#)
- [Configuration Example for EIGRP, page 28-19](#)
- [Feature History for EIGRP, page 28-20](#)

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 not 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.

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 28-4](#)

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 <i>as-num</i>	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.
	Example: ciscoasa(config)# router eigrp 2	The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	network <i>ip-addr</i> [<i>mask</i>]	Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
	Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0	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 28-7 .

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 <i>as-num</i> Example: ciscoasa(config)# router eigrp 2	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	network <i>ip-addr</i> [<i>mask</i>] Example: ciscoasa(config)# router eigrp 2 ciscoasa(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 28-8.
Step 3	eigrp stub { <i>receive-only</i> [<i>connected</i>] [<i>redistributed</i>] [<i>static</i>] [<i>summary</i>]} Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(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.

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 28-6](#)
- [Configuring Interfaces for EIGRP, page 28-7](#)
- [Configuring the Summary Aggregate Addresses on Interfaces, page 28-9](#)
- [Changing the Interface Delay Value, page 28-10](#)
- [Enabling EIGRP Authentication on an Interface, page 28-10](#)
- [Defining an EIGRP Neighbor, page 28-12](#)

- [Redistributing Routes Into EIGRP, page 28-12](#)
- [Filtering Networks in EIGRP, page 28-14](#)
- [Customizing the EIGRP Hello Interval and Hold Time, page 28-15](#)
- [Disabling Automatic Route Summarization, page 28-16](#)
- [Configuring Default Information in EIGRP, page 28-16](#)
- [Disabling EIGRP Split Horizon, page 28-17](#)
- [Restarting the EIGRP Process, page 28-18](#)

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	<code>router eigrp as-num</code>	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.
	Example: ciscoasa(config)# router eigrp 2	The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	<code>network ip-addr [mask]</code>	Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.
	Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0	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 28-8 .

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 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	<pre>router eigrp as-num</pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2</pre></p>	<p>Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.</p> <p>The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.</p>
Step 2	<pre>ciscoasa(config-router)# network ip-addr [mask]</pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0</pre></p>	<p>Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.</p> <p>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.</p> <p>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 28-6.</p>
Step 3	<p>(Optional) Do one of the following to customize an interface to participate in EIGRP routing:</p> <pre>no default-information {in out WORD}</pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(config-router)# no default-information {in out WORD}</pre></p> <pre>authentication mode eigrp as-num md5</pre> <p>Example: <pre>ciscoasa(config)# authentication mode eigrp 2 md5</pre></p>	<p>Allows you to control the sending or receiving of candidate default route information.</p> <p>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.</p> <p>See the “Configuring Default Information in EIGRP” section on page 28-16 for more information on this particular option.</p> <p>Enables MD5 authentication of EIGRP packets.</p> <p>The <i>as-num</i> 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:</p> <pre>% System(100) specified does not exist</pre> <p>See the “Enabling EIGRP Authentication on an Interface” section on page 28-10 for more information on this particular option.</p>

Command	Purpose
<p><code>delay value</code></p> <p>Example: <code>ciscoasa(config-if)# delay 200</code></p>	<p>The <i>value</i> argument entered is in tens of microseconds. To set the delay for 2000 microseconds, you enter a <i>value</i> of 200.</p> <p>To view the delay value assigned to an interface, use the show interface command.</p> <p>See the “Changing the Interface Delay Value” section on page 28-10 for more information on this particular option.</p>
<p><code>hello-interval eigrp as-num seconds</code></p> <p>Example: <code>ciscoasa(config)# hello-interval eigrp 2 60</code></p>	<p>Allows you to change the hello interval. See the “Customizing the EIGRP Hello Interval and Hold Time” section on page 28-15 for more information on this particular option.</p>
<p><code>hold-time eigrp as-num seconds</code></p> <p>Example: <code>ciscoasa(config)# hold-time eigrp 2 60</code></p>	<p>Allows you to change the hold time. See the “Customizing the EIGRP Hello Interval and Hold Time” section on page 28-15 for more information on this particular option.</p>

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	<p><code>router eigrp as-num</code></p> <p>Example: <code>ciscoasa(config)# router eigrp 2</code></p>	<p>Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.</p> <p>The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.</p>

	Command	Purpose
Step 2	<pre>ciscoasa(config-router)# network ip-addr [mask] Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0</pre>	<p>Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.</p> <p>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.</p> <p>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 28-6.</p>
Step 3	<pre>passive-interface {default if-name} Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(config-router)# passive-interface {default}</pre>	<p>Prevents an interface from sending or receiving EIGRP routing message.</p> <p>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.</p>

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	<pre>interface phy_if Example: ciscoasa(config)# interface inside</pre>	<p>Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.</p>
Step 2	<pre>summary-address eigrp as-num address mask [distance] Example: ciscoasa(config-if)# summary-address eigrp 2 address mask [20]</pre>	<p>Creates the summary address.</p> <p>By default, EIGRP summary addresses that you define have an administrative distance of 5. You can change this value by specifying the optional <i>distance</i> argument in the summary-address command.</p>

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	<code>interface phy_if</code>	Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.
	Example: <code>ciscoasa(config)# interface inside</code>	
Step 2	<code>delay value</code>	The <i>value</i> argument entered is in tens of microseconds. To set the delay for 2000 microseconds, you enter a <i>value</i> of 200.
	Example: <code>ciscoasa(config-if)# delay 200</code>	To view the delay value assigned to an interface, use the show interface command.

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	<p>router eigrp <i>as-num</i></p> <p>Example: hostname(config)# router eigrp 2</p>	<p>Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.</p> <p>The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.</p>
Step 2	<p>network <i>ip-addr [mask]</i></p> <p>Example: hostname(config)# router eigrp 2 hostname(config-router)# network 10.0.0.0 255.0.0.0</p>	<p>Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.</p> <p>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.</p> <p>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 28-3.</p>
Step 3	<p>interface <i>phy_if</i></p> <p>Example: hostname(config)# interface inside</p>	<p>Enters interface configuration mode for the interface on which you are configuring EIGRP message authentication.</p>
Step 4	<p>authentication mode eigrp <i>as-num md5</i></p> <p>Example: hostname(config)# authentication mode eigrp 2 md5</p>	<p>Enables MD5 authentication of EIGRP packets.</p> <p>The <i>as-num</i> 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:</p> <pre>% Asystem(100) specified does not exist</pre>
Step 5	<p>authentication key eigrp <i>as-num key key-id key-id</i></p> <p>Example: hostname(config)# authentication key eigrp 2 cisco key-id 200</p>	<p>Configures the key used by the MD5 algorithm.</p> <p>The <i>as-num</i> 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:</p> <pre>% Asystem(100) specified does not exist%</pre> <p>The <i>key</i> argument can include up to 16 characters, including alphabets, numbers and special characters.</p> <p> Note White spaces are not allowed, in the key argument.</p> <p>The <i>key-id</i> argument is a number that can range from 0 to 255.</p>

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	<code>router eigrp as-num</code>	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.
	Example: <code>ciscoasa(config)# router eigrp 2</code>	The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	<code>neighbor ip-addr interface if_name</code>	Defines the static neighbor.
	Example: <code>ciscoasa(config)# router eigrp 2</code> <code>ciscoasa(config-router)# neighbor 10.0.0.0</code> <code>interface interface1</code>	The <i>ip-addr</i> argument is the IP address of the neighbor. The <i>if_name</i> 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.

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	<pre>router eigrp <i>as-num</i></pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2</pre></p>	<p>Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.</p> <p>The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.</p>
Step 2	<pre>default-metric <i>bandwidth delay reliability loading mtu</i></pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# default-metric bandwidth delay reliability loading mtu</pre></p>	<p>(Optional) Specifies the default metrics that should be applied to routes redistributed into the EIGRP routing process.</p> <p>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.</p>
Step 3	<p>Do one of the following to redistribute the selected route type into the EIGRP routing process:</p> <pre>redistribute connected [<i>metric bandwidth delay reliability loading mtu</i>] [<i>route-map map_name</i>]</pre> <p>Example: <pre>ciscoasa(config-router): redistribute connected [<i>metric bandwidth delay reliability loading mtu</i>] [<i>route-map map_name</i>]</pre></p> <pre>redistribute static [<i>metric bandwidth delay reliability loading mtu</i>] [<i>route-map map_name</i>]</pre> <p>Example: <pre>ciscoasa(config-router): redistribute static [<i>metric bandwidth delay reliability loading mtu</i>] [<i>route-map map_name</i>]</pre></p>	<p>Redistributes connected routes into the EIGRP routing process.</p> <p>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.</p> <p>Redistributes static routes into the EIGRP routing process.</p>

Command	Purpose
<pre>redistribute ospf pid [match {internal external [1 2] nssa-external [1 2]}] [metric bandwidth delay reliability loading mtu] [route-map map_name]</pre> <p>Example:</p> <pre>ciscoasa(config-router): redistribute ospf pid [match {internal external [1 2] nssa-external [1 2]}] [metric bandwidth delay reliability loading mtu] [route-map map_name]</pre>	Redistributes routes from an OSPF routing process into the EIGRP routing process.
<pre>redistribute rip [metric bandwidth delay reliability load mtu] [route-map map_name]</pre> <p>Example:</p> <pre>(config-router): redistribute rip [metric bandwidth delay reliability load mtu] [route-map map_name]</pre>	Redistributes routes from a RIP routing process into the EIGRP routing process.

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	<pre>router eigrp as-num</pre> <p>Example:</p> <pre>ciscoasa(config)# router eigrp 2</pre>	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	<pre>ciscoasa(config-router)# network ip-addr [mask]</pre> <p>Example:</p> <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0</pre>	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 28-7.

Command	Purpose
Step 3 Do one of the following to filter networks sent or received in EIGRP routing updates:	
<pre>distribute-list <i>acl</i> out [connected ospf rip static interface <i>if_name</i>]</pre> <p>Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(config-router): distribute-list acl out [connected]</p>	<p>Filters networks sent in EIGRP routing updates.</p> <p>You can specify an interface to apply the filter to only those updates that are sent by that specific interface.</p> <p>You can enter multiple distribute-list commands in your EIGRP router configuration.</p>
<pre>distribute-list <i>acl</i> in [interface <i>if_name</i>]</pre> <p>Example: ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(config-router): distribute-list acl in [interface interface1]</p>	<p>Filters networks received in EIGRP routing updates.</p> <p>You can specify an interface to apply the filter to only those updates that are received by that interface.</p>

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
<p>Step 1 interface <i>phy_if</i></p> <p>Example: ciscoasa(config)# interface inside</p>	<p>Enters interface configuration mode for the interface on which you are configuring the hello interval or advertised hold time.</p>

	Command	Purpose
Step 2	hello-interval eigrp <i>as-num seconds</i> Example: ciscoasa(config)# hello-interval eigrp 2 60	Changes the hello interval.
Step 3	hold-time eigrp <i>as-num seconds</i> Example: ciscoasa(config)# hold-time eigrp 2 60	Changes the hold time.

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 <i>as-num</i> Example: ciscoasa(config)# router eigrp 2	Creates an EIGRP routing process and enters router configuration mode for this EIGRP process. The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.
Step 2	no auto-summary Example: ciscoasa(config-router)# no auto-summary	You cannot configure this value. Automatic summary addresses have an administrative distance of 5.

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	<pre>router eigrp <i>as-num</i></pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2</pre></p>	<p>Creates an EIGRP routing process and enters router configuration mode for this EIGRP process.</p> <p>The <i>as-num</i> argument is the autonomous system number of the EIGRP routing process.</p>
Step 2	<pre>ciscoasa(config-router)# network <i>ip-addr</i> [<i>mask</i>]</pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0</pre></p>	<p>Configures the interfaces and networks that participate in EIGRP routing. You can configure one or more network statements with this command.</p> <p>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.</p> <p>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 28-7.</p>
Step 3	<pre>no default-information {in out WORD}</pre> <p>Example: <pre>ciscoasa(config)# router eigrp 2 ciscoasa(config-router)# network 10.0.0.0 255.0.0.0 ciscoasa(config-router)# no default-information {in out WORD}</pre></p>	<p>Controls the sending or receiving of candidate default route information.</p> <p>Entering the no default-information in command causes the candidate default route bit to be blocked on received routes.</p> <p>Entering the no default-information out command disables the setting of the default route bit in advertised routes.</p>

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	<code>interface phy_if</code> Example: <code>ciscoasa(config)# interface phy_if</code>	Enters interface configuration mode for the interface on which you are changing the delay value used by EIGRP.
Step 2	<code>no split-horizon eigrp as-number</code> Example: <code>ciscoasa(config-if)# no split-horizon eigrp 2</code>	Disables the split horizon.

Restarting the EIGRP Process

To restart an EIGRP process or clear redistribution or counters, enter the following command:

Command	Purpose
<code>clear eigrp pid {1-65535 neighbors topology events}</code> Example: <code>ciscoasa(config)# clear eigrp pid 10 neighbors</code>	Restarts an EIGRP process or clears redistribution or counters.

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 commands:

Command	Purpose
Monitoring EIGRP Routing	
<code>router-id</code>	Displays the router-id for this EIGRP process.
<code>show eigrp [as-number] events [{start end} type]</code>	Displays the EIGRP event log.
<code>show eigrp [as-number] interfaces [if-name] [detail]</code>	Displays the interfaces participating in EIGRP routing.
<code>show eigrp [as-number] neighbors [detail static] [if-name]</code>	Displays the EIGRP neighbor table.

Command (continued)	Purpose (continued)
<code>show eigrp [as-number] topology [ip-addr [mask] active all-links pending summary zero-successors]</code>	Displays the EIGRP topology table.
<code>show eigrp [as-number] traffic</code>	Displays EIGRP traffic statistics.
<code>show mfib cluster</code>	Displays MFIB information in terms of forwarding entries and interfaces.
<code>show route cluster</code>	Displays additional route synchronization details for clustering.
Disabling EIGRP Logging Messages	
<code>no eigrp log-neighbor-changes</code>	Disables the logging of neighbor change messages. Enter this command in router configuration mode for the EIGRP routing process.
<code>no eigrp log-neighbor-warnings</code>	Disables the logging of neighbor warning messages.

**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:

```
ciscoasa(config)# router eigrp 2
ciscoasa(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:

```
ciscoasa(config-router)# passive-interface {default}
```

Step 3 To define an EIGRP neighbor, enter the following command:

```
ciscoasa(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:

```
ciscoasa(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:

```
ciscoasa(config-router)# exit
ciscoasa(config)# interface phy_if
ciscoasa(config-if)# delay 200
```

Feature History for EIGRP

Table 28-1 lists each feature change and the platform release in which it was implemented.

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 .
Dynamic Routing in Multiple Context Mode	9.0(1)	EIGRP routing is supported in multiple context mode.
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 29-4](#)
- [Monitoring RIP, page 29-11](#)
- [Configuration Example for RIP, page 29-11](#)
- [Feature History for RIP, page 29-12](#)

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 29-4](#) 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 25-4](#) 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
<code>router rip</code>	Starts the RIP routing process and places you in router configuration mode.
Example: <code>ciscoasa(config)# router rip</code>	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.

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 29-6](#)
- [Configuring the RIP Send and Receive Version on an Interface, page 29-6](#)
- [Configuring Route Summarization, page 29-7](#)
- [Filtering Networks in RIP, page 29-8](#)
- [Redistributing Routes into the RIP Routing Process, page 29-8](#)
- [Enabling RIP Authentication, page 29-9](#)

- [Restarting the RIP Process, page 29-10](#)

Configuring the RIP Version

To specify the version of RIP used by the ASA, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<code>router rip</code> Example: <code>ciscoasa(config)# router rip</code>	Starts the RIP routing process and places you in router configuration mode.
Step 2	<code>network network_address</code> Example: <code>ciscoasa(config)# router rip</code> <code>ciscoasa(config-router)# network 10.0.0.0</code>	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.
Step 3	Enter one of the following numbers to customize an interface to participate in RIP routing: <code>version [1 2]</code> Example: <code>ciscoasa(config-router)# version [1]</code>	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.

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: ciscoasa(config)# router rip	Starts the RIP routing process and places you in router configuration mode.
Step 2	network network_address Example: ciscoasa(config)# router rip ciscoasa(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: ciscoasa(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.

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: ciscoasa(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.	

Command	Purpose
rip send version {[1] [2]} Example: ciscoasa(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: ciscoasa(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.

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: ciscoasa(config)# router rip	Enables the RIP routing process and places you in router configuration mode.
Step 2	no auto-summarize Example: ciscoasa(config-router) :# no auto-summarize	Disables automatic route summarization.

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	<code>router rip</code>	Enables the RIP routing process and places you in router configuration mode.
	Example: <code>ciscoasa(config)# router rip</code>	
Step 2	<code>distribute-list acl in [interface if_name]</code> <code>distribute-list acl out [connected eigrp interface if_name ospf rip static]</code>	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.
	Example: <code>ciscoasa(config-router)# distribute-list acl2 in [interface interface1]</code> <code>ciscoasa(config-router)# distribute-list acl3 out [connected]</code>	

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
<p>redistribute connected [metric <i>metric-value</i> transparent] [route-map <i>route-map-name</i>]</p> <p>Example: <pre>ciscoasa(config-router): # redistribute connected [metric metric-value transparent] [route-map route-map-name]</pre></p>	<p>Redistributes connected routes into the RIP routing process.</p> <p>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.</p>
<p>redistribute static [metric {<i>metric_value</i> transparent}] [route-map <i>map_name</i>]</p> <p>Example: <pre>ciscoasa(config-router):# redistribute static [metric {metric_value transparent}] [route-map map_name]</pre></p>	<p>Redistributes static routes into the EIGRP routing process.</p>
<p>redistribute ospf <i>pid</i> [match {internal external [1 2] nssa-external [1 2]}] [metric {<i>metric_value</i> transparent}] [route-map <i>map_name</i>]</p> <p>Example: <pre>ciscoasa(config-router):# redistribute ospf pid [match {internal external [1 2] nssa-external [1 2]}] [metric {metric_value transparent}] [route-map map_name]</pre></p>	<p>Redistributes routes from an OSPF routing process into the RIP routing process.</p>
<p>redistribute eigrp <i>as-num</i> [metric {<i>metric_value</i> transparent}] [route-map <i>map_name</i>]</p> <p>Example: <pre>ciscoasa(config-router):# redistribute eigrp as-num [metric {metric_value transparent}] [route-map map_name]</pre></p>	<p>Redistributes routes from an EIGRP routing process into the RIP routing process.</p>

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	<code>router rip as-num</code> Example: ciscoasa(config)# router rip 2	Creates the RIP routing process and enters router configuration mode for this RIP process. The <i>as-num</i> argument is the autonomous system number of the RIP routing process.
Step 2	<code>interface phy_if</code> Example: ciscoasa(config)# interface phy_if	Enters interface configuration mode for the interface on which you are configuring RIP message authentication.
Step 3	<code>rip authentication mode {text md5}</code> Example: ciscoasa(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	<code>rip authentication key key key-id key-id</code> Example: ciscoasa(config-if)# rip authentication key cisco key-id 200	Configures the authentication key used by the MD5 algorithm. The <i>key</i> argument can include up to 16 characters. The <i>key-id</i> argument is a number from 0 to 255.

Restarting the RIP Process

To remove the entire RIP configuration, enter the following command:

Command	Purpose
<code>clear rip pid {process redistribution counters [neighbor [neighbor-interface] [neighbor-id]]}</code> Example: ciscoasa(config)# clear rip	Removes the entire RIP configuration that you have enabled. After the configuration is cleared, you must reconfigure RIP again using the router rip command.

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	
<code>show rip database</code>	Display the contents of the RIP routing database.
<code>show running-config router rip</code>	Displays the RIP commands.
<code>show route cluster</code>	Displays additional route synchronization details for clustering.
Debugging RIP	
<code>debug rip events</code>	Displays RIP processing events.
<code>debug rip database</code>	Displays RIP database events.
<code>debug route cluster</code>	Enables RIB table replication trace messages to determine if the RIB is correctly synchronized to the slave units in clustering.

Configuration Example for RIP

The following example shows how to enable and configure RIP with various optional processes:

```
ciscoasa(config)# router rip 2
ciscoasa(config-router)# default-information originate
ciscoasa(config-router)# version [1]
ciscoasa(config-router)# network 225.25.25.225
ciscoasa(config-router)# passive-interface [default]
ciscoasa(config-router)# redistribute connected [metric bandwidth delay reliability loading mtu] [route-map map_name]
```

Feature History for RIP

Table 29-1 lists each feature change and the platform release in which it was implemented.

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.
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 30-3](#)
- [Customizing Multicast Routing, page 30-4](#)
- [Configuration Example for Multicast Routing, page 30-15](#)
- [Additional References, page 30-15](#)
- [Feature History for Multicast Routing, page 30-16](#)

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 30-2](#)

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the ASA acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the ASA forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the ASA cannot be configured for PIM.

The ASA supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees rooted at a single Rendezvous Point per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bi-directional PIM is a variant of PIM-SM that builds bi-directional shared trees connecting multicast sources and receivers. Bi-directional trees are built using a DF election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point, and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during Rendezvous Point discovery and provides a default route to the Rendezvous Point.

**Note**

If the ASA is the PIM 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.

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 8, “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 command:

Command	Purpose
<code>multicast-routing</code>	Enables multicast routing.
Example: <code>ciscoasa(config)# multicast-routing</code>	The number of entries in the multicast routing tables are limited by the amount of RAM on the ASA.

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 30-4](#)
- [Configuring a Static Multicast Route, page 30-5](#)
- [Configuring IGMP Features, page 30-5](#)
- [Configuring PIM Features, page 30-10](#)
- [Configuring a Bidirectional Neighbor Filter, page 30-13](#)
- [Configuring a Multicast Boundary, page 30-14](#)

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 area:

Command	Purpose
<pre>igmp forward interface <i>if_name</i></pre> <p>Example: <pre>ciscoasa(config-if)# igmp forward interface <i>interface1</i></pre></p>	Configures stub multicast routing and forwards IGMP messages.

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 commands:

Command	Purpose
<pre>mroute <i>src_ip src_mask</i> {<i>input_if_name</i> <i>rpf_neighbor</i>} [<i>distance</i>]</pre> <p>Example: <pre>ciscoasa(config)# mroute <i>src_ip src_mask</i> {<i>input_if_name</i> <i>rpf_neighbor</i>} [<i>distance</i>]</pre></p>	Configures a static multicast route.
<pre>mroute <i>src_ip src_mask input_if_name</i> [dense <i>output_if_name</i>] [<i>distance</i>]</pre> <p>Example: <pre>ciscoasa(config)# mroute <i>src_ip src_mask</i> <i>input_if_name</i> [dense <i>output_if_name</i>] [<i>distance</i>]</pre></p>	Configures a static multicast route for a stub area. The dense <i>output_if_name</i> keyword and argument pair is only supported for stub multicast routing.

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 30-6](#)
- [Configuring IGMP Group Membership, page 30-7](#)
- [Configuring a Statically Joined IGMP Group, page 30-7](#)
- [Controlling Access to Multicast Groups, page 30-8](#)
- [Limiting the Number of IGMP States on an Interface, page 30-8](#)
- [Modifying the Query Messages to Multicast Groups, page 30-8](#)
- [Changing the IGMP Version, page 30-9](#)

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 command:

Command	Purpose
<code>no igmp</code>	Disables IGMP on an interface. To reenable IGMP on an interface, use the igmp command.
Example: <code>ciscoasa(config-if)# no igmp</code>	

**Note**

Only the **no igmp** command appears in the interface configuration.

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 30-7](#).

To have the ASA join a multicast group, enter the following command:

Command	Purpose
igmp join-group <i>group-address</i> Example: ciscoasa(config-if)# igmp join-group <i>mcast-group</i>	Configures the ASA to be a member of a multicast group. The <i>group-address</i> argument is the IP address of the group.

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.

To configure a statically joined multicast group on an interface, enter the following command:

Command	Purpose
igmp static-group Example: ciscoasa(config-if)# igmp static-group <i>group-address</i>	Configures the ASA statically to join a multicast group on an interface. The <i>group-address</i> argument is the IP address of the group.

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:	
	<pre>access-list <i>name</i> standard [permit deny] <i>ip_addr mask</i></pre> <p>Example: <pre>ciscoasa(config)# access-list <i>acl1</i> standard permit 192.52.662.25</pre></p>	<p>Creates a standard ACL for the multicast traffic.</p> <p>You can create more than one entry for a single ACL. You can use extended or standard ACLs.</p> <p>The <i>ip_addr mask</i> argument is the IP address of the multicast group being permitted or denied.</p>
	<pre>access-list <i>name</i> extended [permit deny] <i>protocol src_ip_addr src_mask dst_ip_addr dst_mask</i></pre> <p>Example: <pre>ciscoasa(config)# access-list <i>acl2</i> extended permit <i>protocol</i> <i>src_ip_addr</i> <i>src_mask dst_ip_addr dst_mask</i></pre></p>	<p>Creates an extended ACL.</p> <p>The <i>dst_ip_addr</i> argument is the IP address of the multicast group being permitted or denied.</p>
Step 2	<pre>igmp access-group <i>acl</i></pre> <p>Example: <pre>ciscoasa(config-if)# igmp access-group <i>acl</i></pre></p>	<p>Applies the ACL to an interface.</p> <p>The <i>acl</i> argument is the name of a standard or extended IP ACL.</p>

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 command:

Command	Purpose
<pre>igmp limit <i>number</i></pre> <p>Example: <pre>ciscoasa(config-if)# igmp limit 50</pre></p>	<p>Limits the number of IGMP states on an interface.</p> <p>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.</p>

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	<code>igmp query-interval seconds</code> Example: <code>ciscoasa(config-if)# igmp query-interval 30</code>	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	<code>igmp query-timeout seconds</code> Example: <code>ciscoasa(config-if)# igmp query-timeout 30</code>	Changes the timeout value of the query. Valid values range from 0 to 500; 225 is the default value.
Step 3	<code>igmp query-max-response-time seconds</code> Example: <code>ciscoasa(config-if)# igmp query-max-response-time 30</code>	Changes the maximum query response time.

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 command:

Command	Purpose
<code>igmp version {1 2}</code>	Controls the version of IGMP that you want to run on the interface.
Example: <code>ciscoasa(config-if)# igmp version 2</code>	

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 30-10](#)
- [Configuring a Static Rendezvous Point Address, page 30-11](#)
- [Configuring the Designated Router Priority, page 30-11](#)
- [Configuring and Filtering PIM Register Messages, page 30-12](#)
- [Configuring PIM Message Intervals, page 30-12](#)
- [Filtering PIM Neighbors, page 30-12](#)

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	<code>pim</code> Example: <code>ciscoasa(config-if)# pim</code>	Enables or reenables PIM on a specific interface.
Step 2	<code>no pim</code> Example: <code>ciscoasa(config-if)# no pim</code>	Disables PIM on a specific interface.



Note

Only the **no pim** command appears in the interface configuration.

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 RP, enter the following command:

Command	Purpose
<p>pim rp-address <i>ip_address</i> [<i>acl</i>] [bidir]</p> <p>Example: <pre>ciscoasa(config)# pim rp-address 10.86.75.23 [acl1] [bidir]</pre></p>	<p>Enables or reenables PIM on a specific interface.</p> <p>The <i>ip_address</i> argument is the unicast IP address of the router assigned to be a PIM RP.</p> <p>The <i>acl</i> 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.</p> <p>Excluding the bidir keyword causes the groups to operate in PIM sparse mode.</p>



Note

The ASA always advertises the bidirectional capability in the PIM hello messages, regardless of the actual bidirectional configuration.

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 command:

Command	Purpose
<p>pim dr-priority <i>num</i></p> <p>Example: <pre>ciscoasa(config-if)# pim dr-priority 500</pre></p>	<p>Changes the designated router priority.</p> <p>The <i>num</i> argument can be any number ranging from 1 to 4294967294.</p>

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 command:

Command	Purpose
<code>pim accept-register {list <i>acl</i> route-map <i>map-name</i>}</code>	Configures the ASA to filter PIM register messages.
Example: <code>ciscoasa(config)# pim accept-register {list <i>acl1</i> route-map <i>map2</i>}</code>	In the example, the ASA filters PIM register messages <i>acl1</i> and route map <i>map2</i> .

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	<code>pim hello-interval <i>seconds</i></code> Example: <code>ciscoasa(config-if)# pim hello-interval 60</code>	Sends router query messages. Valid values for the <i>seconds</i> argument range from 1 to 3600 seconds.
Step 2	<code>pim join-prune-interval <i>seconds</i></code> Example: <code>ciscoasa(config-if)# pim join-prune-interval 60</code>	Changes the amount of time (in seconds) that the ASA sends PIM join or prune messages. Valid values for the <i>seconds</i> argument range from 10 to 600 seconds.

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	<pre>access-list pim_nbr deny router-IP_addr PIM neighbor</pre> <p>Example: <pre>ciscoasa(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255</pre></p>	<p>Uses a standard ACL to define the routers that you want to have participate in PIM.</p> <p>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.</p>
Step 2	<pre>pim neighbor-filter pim_nbr</pre> <p>Example: <pre>ciscoasa(config)# interface GigabitEthernet0/3 ciscoasa(config-if)# pim neighbor-filter pim_nbr</pre></p>	<p>Filters neighbor routers.</p> <p>In the example, the 10.1.1.1 router is prevented from becoming a PIM neighbor on interface GigabitEthernet0/3.</p>

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	<pre>access-list pim_nbr deny router-IP_addr PIM neighbor</pre> <p>Example: <pre>ciscoasa(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255</pre></p>	<p>Uses a standard ACL to define the routers that you want to have participate in PIM.</p> <p>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.</p>
Step 2	<pre>pim bidirectional-neighbor-filter pim_nbr</pre> <p>Example: <pre>ciscoasa(config)# interface GigabitEthernet0/3 ciscoasa(config-if)# pim bidirectional neighbor-filter pim_nbr</pre></p>	<p>Filters neighbor routers.</p> <p>In the example, the 10.1.1.1 router is prevented from becoming a PIM bidirectional neighbor on interface GigabitEthernet0/3.</p>

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 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 command:

Command	Purpose
<pre>multicast boundary acl [filter-autorp]</pre> <p>Example: <pre>ciscoasa(config-if)# multicast boundary acl1 [filter-autorp]</pre></p>	<p>Configures a multicast boundary.</p>

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:

```
ciscoasa(config)# multicast-routing
```

Step 2 Configure a static multicast route:

```
ciscoasa(config)# mroute src_ip src_mask {input_if_name | rpf_neighbor} [distance]  
ciscoasa(config)# exit
```

Step 3 Configure the ASA to be a member of a multicast group:

```
ciscoasa(config)# interface  
ciscoasa(config-if)# igmp join-group group-address
```

Additional References

For additional information related to routing, see the following sections:

- [Related Documents, page 30-16](#)
- [RFCs, page 30-16](#)

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.

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.
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 31-4](#)
- [Prerequisites for IPv6 Neighbor Discovery, page 31-4](#)
- [Guidelines and Limitations, page 31-4](#)
- [Default Settings for IPv6 Neighbor Discovery, page 31-6](#)
- [Configuring IPv6 Neighbor Discovery, page 31-6](#)
- [Monitoring IPv6 Neighbor Discovery, page 31-14](#)
- [Additional References, page 31-14](#)
- [Feature History for IPv6 Neighbor Discovery, page 31-15](#)

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 cached 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 31-2](#)
- [Duplicate Address Detection, page 31-2](#)
- [Router Advertisement Messages, page 31-3](#)
- [Static IPv6 Neighbors, page 31-4](#)

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 (ICMPv6 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.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verify the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-nodes multicast address.

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 on 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.

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 11-12.

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 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
<i>value</i> for the neighbor solicitation transmission message interval	1000 seconds between neighbor solicitation transmissions.
<i>value</i> for the neighbor reachable time	The default is 0.
<i>value</i> for the router advertisement transmission interval	The default is 200 seconds.
<i>value</i> for the router lifetime	The default is 1800 seconds.
<i>value</i> 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 31-6](#)
- [Configuring the Neighbor Solicitation Message Interval, page 31-7](#)
- [Configuring the Neighbor Reachable Time, page 31-8](#)
- [Configuring the Router Advertisement Transmission Interval, page 31-8](#)
- [Configuring the Router Lifetime Value, page 31-9](#)
- [Configuring DAD Settings, page 31-9](#)
- [Suppressing Router Advertisement Messages, page 31-10](#)
- [Configuring Address Config Flags for IPv6 DHCP Relay, page 31-11](#)
- [Configuring the IPv6 Prefix in Router Advertisements, page 31-12](#)
- [Configuring a Static IPv6 Neighbor, page 31-13](#)

Entering Interface Configuration Mode

Configure neighbor discovery settings per interface. To enter interface configuration mode, perform the following steps.

Detailed Steps

Command	Purpose
<code>interface name</code>	Enters interface configuration mode.
Example: <pre>hostname(config)# interface gigabitethernet 0/0 hostname(config-if)#</pre>	

Configuring the Neighbor Solicitation Message Interval

To configure the interval between IPv6 neighbor solicitation retransmissions on an interface, enter the following command.

Detailed Steps

Command	Purpose
<code>ipv6 nd ns-interval value</code>	Sets the interval between IPv6 neighbor solicitation retransmissions on an interface.
Example: <pre>hostname (config-if)# ipv6 nd ns-interval 9000</pre>	Valid values for the value argument range from 1000 to 3600000 milliseconds. This information is also sent in router advertisement messages.

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 command.

Detailed Steps

Command	Purpose
<code>ipv6 nd reachable-time value</code>	Sets the amount of time that a remote IPv6 node is reachable.
Example: hostname (config-if)# ipv6 nd reachable-time 1700000	Valid values for the <i>value</i> 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.

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 command.

Detailed Steps

Command	Purpose
<code>ipv6 nd ra-interval [msec] value</code>	Sets the interval between IPv6 router advertisement transmissions.
Example: hostname (config-if)# ipv6 nd ra-interval 201	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 <i>value</i> 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 31-9 . To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value.

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 command.

Detailed Steps

Command	Purpose
<code>ipv6 nd ra-lifetime [msec] value</code>	Specifies the length of time that nodes on the local link should consider the ASA as the default router on the link.
Example: hostname (config-if)# ipv6 nd ra-lifetime 2000	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 <i>value</i> argument range from 0 to 9000 seconds. Entering 0 indicates that the ASA should not be considered a default router on the selected interface.

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 command.

Detailed Steps

Command	Purpose
<code>ipv6 nd dad attempts value</code>	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.
Example: hostname (config-if)# ipv6 nd dad attempts 20	Valid values for the <i>value</i> argument range from 0 to 600. A zero value disables DAD processing on the specified interface.

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 command.

Detailed Steps

Command	Purpose
ipv6 nd suppress-ra	Suppresses IPv6 router advertisement transmissions on a LAN interface.
Example: hostname (config-if)# ipv6 nd suppress-ra	Entering this command causes the ASA to appear as a regular IPv6 neighbor on the link and not as an IPv6 router.

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
```


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.

Configuring the IPv6 Prefix in Router Advertisements

To configure the which IPv6 prefixes are included in IPv6 router advertisements, enter the following command.

Detailed Steps

Command	Purpose
<pre> ipv6 nd prefix <i>ipv6-prefix/prefix-length</i> default [[<i>valid-lifetime</i> <i>preferred-lifetime</i>] [at <i>valid-date</i> <i>preferred-date</i>] infinite no-advertise off-link no-autoconfig] </pre> <p>Example:</p> <pre> hostname (config-if)# ipv6 nd prefix 2001:DB8::/32 1000 900 </pre>	<p>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.</p> <p>The at <i>valid-date preferred-date</i> 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 <i>date-valid-expire month-valid-expire hh:mm-valid-expire date-prefer-expire month-prefer-expire hh:mm-prefer-expire</i>.</p> <p>The default keyword indicates that default values are used.</p> <p>The optional infinite keyword specifies that the valid lifetime does not expire.</p> <p>The <i>ipv6-prefix</i> 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.</p> <p>The optional no-advertise keyword indicates to hosts on the local link that the specified prefix is not to be used for IPv6 autoconfiguration.</p> <p>The optional no-autoconfig keyword indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.</p> <p>The optional off-link keyword indicates that the specified prefix is not used for on-link determination.</p> <p>The <i>preferred-lifetime</i> 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).</p> <p>The <i>prefix-length</i> 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.</p> <p>The <i>valid-lifetime</i> 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).</p>

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.

Detailed Steps

Command	Purpose
<pre>ipv6 neighbor ipv6_address if_name mac_address</pre> <p>Example:</p> <pre>hostname(config-if)# ipv6 neighbor 3001:1::45A inside 002.7D1A.9472</pre>	<p>Configures a static entry in the IPv6 neighbor discovery cache.</p> <p>The <i>ipv6_address</i> argument is the link-local IPv6 address of the neighbor, the <i>if_name</i> argument is the interface through which the neighbor is available, and the <i>mac_address</i> argument is the MAC address of the neighbor interface.</p>

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
```

Monitoring IPv6 Neighbor Discovery

To monitor IPv6 neighbor discovery parameters, enter the following command:

Command	Purpose
<code>show ipv6 interface</code>	<p>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:</p> <ul style="list-style-type: none"> • 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 31-15](#)
- [RFCs for IPv6 Prefixes and Documentation, page 31-15](#)

Related Documents for IPv6 Prefixes

Related Topic	Document Title
ipv6 commands	<i>command reference</i>

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 <i>ipv6-prefix</i> 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.

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 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 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 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 41, “Configuring Management Access.”](#)
 - Serial console
 - ASDM using HTTPS
 - VPN management access
- The **enable** command. For more information, see [Chapter 41, “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 Chapter 7, “Configuring AAA Rules for Network Access” of the firewall configuration guide.
- VPN access. For more information, see Chapter 6, “Configuring Remote Access IPsec VPNs,” Chapter 8, “Configuring Easy VPN Services on the ASA 5505,” Chapter 10, “Configuring LAN-to-LAN IPsec VPNs,” and Chapter 82, “Introduction to Clientless SSL VPN” of 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 41, “Configuring Management Access.”](#)
- Network access. For more information, see Chapter 7, “Configuring AAA Rules for Network Access” of the firewall configuration guide.
- VPN access. For more information, see Chapter 6, “Configuring Remote Access IPsec VPNs,” Chapter 8, “Configuring Easy VPN Services on the ASA 5505,” Chapter 10, “Configuring LAN-to-LAN IPsec VPNs,” and Chapter 82, “Introduction to Clientless SSL VPN” of 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 Chapter 87, “Configuring Clientless SSL VPN Users,” of the VPN configuration guide.
Identity Firewall	See Chapter 38, “Configuring the Identity Firewall.”
Kerberos	See the “Microsoft Kerberos Constrained Delegation Solution” section on page 84-1 of 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 following chapters of the VPN configuration guide: <ul style="list-style-type: none"> • Chapter 1, “Configuring IPsec and ISAKMP” • Chapter 3, “Setting General VPN Parameters” • Chapter 4, “Configuring Connection Profiles, Group Policies, and Users” • Chapter 6, “Configuring Remote Access IPsec VPNs,” • Chapter 8, “Configuring Easy VPN Services on the ASA 5505” • Chapter 10, “Configuring LAN-to-LAN IPsec VPNs”
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 AAA and 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 33-2](#)
- [Guidelines and Limitations, page 33-3](#)
- [Adding a User Account to the Local Database, page 33-4](#)
- [Monitoring the Local Database, page 33-8](#)
- [Feature History for the Local Database, page 33-9](#)

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.

- 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 41-36.

Adding a User Account to the Local Database

To add a user to the local database, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>username username {nopassword password password} [privilege priv_level]</pre> <p>Example:</p> <pre>ciscoasa(config)# username exampleuser1 privilege 1</pre>	<p>Creates the user account. The username <i>username</i> keyword is a string from 4 to 64 characters long.</p> <p>The password <i>password</i> keyword is a string from 3 to 32 characters long. The privilege <i>level</i> argument sets the privilege level, which ranges from 0 to 15. The default is 2. This privilege level is used with command authorization.</p> <div style="border: 1px solid black; padding: 5px; margin: 10px 0;">  <p>Caution If you do not use command authorization (the aaa authorization console LOCAL command), then the default level 2 allows management access to privileged EXEC mode. If you want to limit access to privileged EXEC mode, either set the privilege level to 0 or 1, or use the service-type command.</p> </div> <p>The nopassword keyword creates a user account with no password.</p> <p>The encrypted keyword indicates that the password is encrypted. When you define a password in the username command, the ASA encrypts it when it saves it to the configuration for security purposes. When you enter the show running-config command, the username command does not show the actual password; it shows the encrypted password followed by the encrypted keyword. For example, if you enter the password “test,” the show running-config output would appear as something similar to the following:</p> <pre>username user1 password DLaUiAX3l78qgoB5c7iVNw== encrypted</pre> <p>The only time you would actually enter the encrypted keyword at the CLI is if you are cutting and pasting a configuration file for use in another ASA, and you are using the same password.</p>

	Command	Purpose
Step 2	<p>username <i>username</i> attributes</p> <p>Example: <pre>ciscoasa(config)# username exampleuser1 attributes</pre></p>	<p>(Optional) Configures username attributes. The <i>username</i> argument is the username that you created in Step 1.</p> <p>By default, VPN users that you add with this command have no attributes or group policy association. You must configure all values explicitly using the username attributes command. For more information, see the Configuring Attributes for Individual Users section in the VPN configuration guide.</p>
Step 3	<p>service-type {admin nas-prompt remote-access}</p> <p>Example: <pre>ciscoasa(config-username)# service-type admin</pre></p>	<p>(Optional) Configures the user level if you configured management authorization using the aaa authorization exec command (see the “Limiting User CLI and ASDM Access with Management Authorization” section on page 41-23). The admin keyword allows full access to any services specified by the aaa authentication console LOCAL commands. The admin keyword is the default.</p> <p>The nas-prompt keyword allows access to the CLI when you configure the aaa authentication {telnet ssh serial} console command, but denies ASDM configuration access if you configure the aaa authentication http console command. ASDM monitoring access is allowed. If you enable authentication with the aaa authentication enable console command, the user cannot access privileged EXEC mode using the enable command (or the login command).</p> <p>The remote-access keyword denies management access. The user cannot use any services specified by the aaa authentication console commands (excluding the serial keyword; serial access is allowed).</p>

	Command	Purpose
Step 4	<pre>ssh authentication {pkf publickey key [hashed]}</pre> <p>Example: ciscoasa(config-username)# ssh authentication pkf</p> <p>Enter an SSH public key formatted file. End with the word "quit" on a line by itself: ---- BEGIN SSH2 PUBLIC KEY ---- Comment: "4096-bit RSA, converted by xxx@xxx from OpenSSH" AAAAB3NzaC1yc2EAAAADAQABAAQADNDUvkgza371B/Q/fljp LAv1BbyAd5PJCjXh/U4LO h1eR/qgIROjpnFaS7Az8/+sjHmq0qXC5TXkzWihvRZbhefyPh PHCi0hIt4oUF2ZbXESA/8 jUT4ehXIUE7FrChffBBtd4d9FkV8A2gwZCDJBxEM26ocbZCS Tx9QC//wt6E/zRcdoqiJG p4ECEdDaM+561+yf73NUigO7wYkqcrzjmI1rZRDLVcqtj8Q9q D3MqsV+PkJGSGiqZwnyI1 QbfYxXHU9wLdWxhUbA/xOjJuZ15TQMa7KLS2u+RtrpQgeTGTF fIh6O+xKh93gwTgzaZTK4 CQ1kuMrRdNRzza0byLeYpTslv6Lv6F6dGtwlqrX5a+w/tV/aw 9WUg/rapekK1oz3tsPTDe p866AFzU+Z7pVR1389iNuNjHQS7IUA2m0cciTuCM2we/tVqMP YJ1+xgKakuHDkBlMS4i8b Wzyd+4EUMDGGZVe0+corKTLWFO1wIUieRkrUaCzjComGYZdZr QT2mXBcSKQNWlSCBpCHsk /r5uTGnKpCNwfl7vd/sRCHyHKsxjsXR15C/5zgHmCTAaGouIq 0Rjo34+61+70PctYXebxM Wwm19e3eH2PudZd+rj1dedfr2/Iris1EBRJWGLoR/N+xsvwVV M1QqwluL4r99CbZF9NghY NRxCQOY/7K77II== ---- END SSH2 PUBLIC KEY ----quit INFO: Import of an SSH public key formatted file SUCCEEDED. ciscoasa(config-username)# </p>	<p>Enables public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key (the pkf keyword) or a Base64 key (the publickey keyword).</p> <p>For a publickey, the <i>key</i> is 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).</p> <p>For a pkf key, you are prompted to paste in a PKF formatted key, up to 4096 bits. Use this format for keys that are too large to paste inline in Base64 format. For example, you can generate a 4096-bit key using ssh keygen, then convert it to PKF, and use the pkf keyword to be prompted for the key.</p> <p>Note You can use the pkf option with failover, but the PKF key is not automatically replicated to the standby system. You must enter the write standby command to synchronize the PKF key.</p> <p>When you view the key on the ASA using the show running-config username command, the key is encrypted using a SHA-256 hash. Even if you entered the key as pkf, the ASA hashes the key, and shows it as a hashed publickey. If you need to copy the key from show output, specify the publickey type with the hashed keyword.</p>
Step 5	<p>(Optional) If you are using this username for VPN authentication, you can configure many VPN attributes for the user. For more information, see the Configuring Attributes for Individual Users section in the VPN configuration guide.</p>	

Examples

The following example assigns a privilege level of 15 to the admin user account:

```
ciscoasa(config)# username admin password password privilege 15
```

The following example creates a user account with no password:

```
ciscoasa(config)# username user34 nopassword
```

The following example enables management authorization, creates a user account with a password, enters username configuration mode, and specifies a **service-type** of **nas-prompt**:

```
ciscoasa(config)# aaa authorization exec authentication-server
ciscoasa(config)# username user1 password gOgeOus
ciscoasa(config)# username user1 attributes
ciscoasa(config-username)# service-type nas-prompt
```

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:

```

jcrichon-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@jcrichon-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:

```

jcrichon-mac:~ john$ cd .ssh
jcrichon-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"
AAAAB3NzaC1yc2EAAAADAQABAAQADNuvkgza371B/Q/fljplAv1BbyAd5PJCjXh/U4LO
hleR/qgIROjpnDas7Az8/+sjHmq0qXC5TXkzWihvRZbhefyPhPHCi0hIt4oUF2ZbXESA/8
jUT4ehXIUE7FrChffBBtd4d9FkV8A2gwZCDJBxEM26ocbZCSTx9QC//wt6E/zRcdqiJG
p4ECEdDaM+56l+yf73NUigO7wYkqcrzjmI1rZRDLVcqtj8Q9qD3MqsV+PkJGSGiqZwnyI1
QbfYxXHU9wLdWxhUba/xOjJuZ15TQMa7KLS2u+RtrpQgeTGTffIh6O+xKh93gwTgzaZTK4
CQ1kuMrRdNRzza0byLeYptSlv6Lv6F6dGtwlqrX5a+w/tV/aw9WUG/rapekKloz3tsPTDe
p866AFzU+Z7pVR1389iNuNJHQs7IUA2mOcciIuCM2we/tVqMPYJl+xgKakuHDkBlMS4i8b
Wzyd+4EUMDGGZVeO+corKTLWF01wIUieRkrUaCzjComGYZdZrQT2mXBcSKQNWlSCBpChsk
/r5uTgnKpCNwFL7vd/sRCHYHksxjsXR15C/5zgHmCTAaGOuIq0Rjo34+6l+70PctYXebxM
Wwm19e3eH2PudZd+rj1dedfr2/IrisLEBRJWGLoR/N+xsvvVVM1Qqw1uL4r99CbZf9NghY
NRxCQOY/7K77IQ==
---- END SSH2 PUBLIC KEY ----
jcrichon-mac:~.ssh john$

```

Step 3 Copy the key to your clipboard.

Step 4 Connect to the ASA CLI, and add the public key to your username:

```

ciscoasa(config)# username test attributes
ciscoasa(config-username)# ssh authentication pkf
Enter an SSH public key formatted file.
End with the word "quit" on a line by itself:
---- BEGIN SSH2 PUBLIC KEY ----
Comment: "4096-bit RSA, converted by ramona@rboersma-mac from OpenSSH"
AAAAB3NzaC1yc2EAAAADAQABAAQADNuvkgza371B/Q/fljplAv1BbyAd5PJCjXh/U4LO
hleR/qgIROjpnDas7Az8/+sjHmq0qXC5TXkzWihvRZbhefyPhPHCi0hIt4oUF2ZbXESA/8
jUT4ehXIUE7FrChffBBtd4d9FkV8A2gwZCDJBxEM26ocbZCSTx9QC//wt6E/zRcdqiJG
p4ECEdDaM+56l+yf73NUigO7wYkqcrzjmI1rZRDLVcqtj8Q9qD3MqsV+PkJGSGiqZwnyI1
QbfYxXHU9wLdWxhUba/xOjJuZ15TQMa7KLS2u+RtrpQgeTGTffIh6O+xKh93gwTgzaZTK4

```

```

CQ1kuMrRdNRzza0byLeYPtSlv6Lv6F6dGtwlqrX5a+w/tV/aw9WUg/rapekKloz3tsPTDe
p866AFzU+Z7pVR1389iNuNJHQS7IUA2m0cciIuCM2we/tVqMPYJl+XgKakuHDkBlMS4i8b
Wzyd+4EUMDGGZVeO+corKTLWFO1wIUieRkrUaCzjComGYZdZrQT2mXBcSKQNW1SCBpCHsk
/r5uTGnKpCNwfL7vd/sRCHYHKsxjsXR15C/5zgHmCTAaGouIq0Rjo34+61+70PctYXebxM
Wwm19e3eH2PudZd+rj1dedfr2/Iris1EBRJWGLoR/N+xsvvVVM1Qqwlul4r99CbZF9NghY
NRxCQOY/7K77IQ==
---- END SSH2 PUBLIC KEY ----
quit
INFO: Import of an SSH public key formatted file completed successfully.

```

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>

```

Monitoring the Local Database

To monitor the local database, enter one of the following commands:

Command	Purpose
<code>show aaa-server</code>	Shows the configured database statistics. To clear the AAA server configuration, enter the clear aaa-server statistics command.
<code>show running-config aaa-server</code>	Shows the AAA server running configuration. To clear AAA server statistics, enter the clear configure aaa-server command.

Feature History for the Local Database

Table 33-1 lists each feature change and the platform release in which it was implemented.

Table 33-1 Feature History for the Local Database

Feature Name	Platform Releases	Feature Information
Local database configuration for AAA	7.0(1)	<p>Describes how to configure the local database for AAA use.</p> <p>We introduced the following commands:</p> <p>username, aaa authorization exec authentication-server, aaa authentication console LOCAL, aaa authorization exec LOCAL, service-type, aaa authentication {telnet ssh serial} console LOCAL, aaa authentication http console LOCAL, aaa authentication enable console LOCAL, show running-config aaa-server, show aaa-server, clear configure aaa-server, clear aaa-server statistics.</p>
Support for SSH public key authentication	9.1(2)	<p>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).</p> <p>We introduced the following commands: ssh authentication.</p> <p><i>Also available in 8.4(4.1); PKF key format support is only in 9.1(2).</i></p>



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](#)
- [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 34-1](#)
- [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: <ul style="list-style-type: none"> • <i>group policy name</i> • <i>OU=group policy name</i> • <i>OU=group policy name;</i>
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: <i>Engineering, Sales</i> 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	0x7ffffff
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 <i>SSL VPN Deployment Guide</i> 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 <i>SSL VPN Deployment Guide</i> 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	0x7ffffff
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	0x7ffffff
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: <ul style="list-style-type: none"> <i>group policy name</i> <i>OU=group policy name</i> <i>OU=group policy name</i>
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: <ul style="list-style-type: none"> .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 41-36](#).

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 34-17](#)

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 34-17.

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

	Command	Purpose
Step 1	aaa-server <i>server_tag</i> protocol radius Example: <pre>ciscoasa(config)# aaa-server servergroup1 protocol radius ciscoasa(config-aaa-server-group)#</pre>	Identifies the server group name and the protocol. When you enter the aaa-server protocol command, you enter aaa-server group configuration mode.
Step 2	merge-dacl { before-avpair after-avpair } Example: <pre>ciscoasa(config)# aaa-server servergroup1 protocol radius ciscoasa(config-aaa-server-group)# merge-dacl before-avpair</pre>	Merges a downloadable ACL with the ACL received in the Cisco AV pair from a RADIUS packet. The default setting is no merge dacl , which specifies that downloadable ACLs will not be merged with Cisco AV pair ACLs. If both an AV pair and a downloadable ACL are received, the AV pair has priority and is used. The before-avpair option specifies that the downloadable ACL entries should be placed before the Cisco AV pair entries. The after-avpair option specifies that the downloadable ACL entries should be placed after the Cisco AV pair entries. This option applies only to VPN connections. For VPN users, ACLs can be in the form of Cisco AV pair ACLs, downloadable ACLs, and an ACL that is configured on the ASA. This option determines whether or not the downloadable ACL and the AV pair ACL are merged, and does not apply to any ACLs configured on the ASA.

	Command	Purpose
Step 3	<p>max-failed-attempts <i>number</i></p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# max-failed-attempts 2</pre></p>	<p>Specifies the maximum number of requests sent to a RADIUS server in the group before trying the next server. The <i>number</i> argument can range from 1 and 5. The default is 3.</p> <p>If you configured a fallback method using the local database (for management access only), and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (by default), so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see the reactivation-mode command in the next step.</p> <p>If you do not have a fallback method, the ASA continues to retry the servers in the group.</p>
Step 4	<p>reactivation-mode {depletion [deadtime <i>minutes</i>] timed}</p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# reactivation-mode deadtime 20</pre></p>	<p>Specifies the method (reactivation policy) by which failed servers in a group are reactivated.</p> <p>The depletion keyword reactivates failed servers only after all of the servers in the group are inactive.</p> <p>The deadtime <i>minutes</i> keyword-argument pair specifies the amount of time in minutes, between 0 and 1440, that elapses between the disabling of the last server in the group and the subsequent reenabling of all servers. The default is 10 minutes.</p> <p>The timed keyword reactivates failed servers after 30 seconds of down time.</p>
Step 5	<p>accounting-mode simultaneous</p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# accounting-mode simultaneous</pre></p>	<p>Sends accounting messages to all servers in the group.</p> <p>To restore the default of sending messages only to the active server, enter the accounting-mode single command.</p>
Step 6	<p>aaa-server <i>server_group</i> [<i>interface_name</i>] host <i>server_ip</i></p> <p>Example: <pre>ciscoasa(config)# aaa-server servergroup1 outside host 10.10.1.1</pre></p>	<p>Identifies the server and the AAA server group to which it belongs.</p> <p>When you enter the aaa-server host command, you enter aaa-server host configuration mode.</p>

Examples

The following example shows how to add one RADIUS group with a single server:

```
ciscoasa(config)# aaa-server AuthOutbound protocol radius
ciscoasa(config-aaa-server-group)# exit.
ciscoasa(config)# aaa-server AuthOutbound (inside) host 10.1.1.3
ciscoasa(config-aaa-server-host)# key RadUauthKey
ciscoasa(config-aaa-server-host)# exit
```

Adding a RADIUS Server to a Group

To add a RADIUS server to a group, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>aaa-server server_group [interface_name] host server_ip</pre> <p>Example: <pre>ciscoasa(config-aaa-server-group)# aaa-server servergroup1 outside host 10.10.1.1</pre></p>	<p>Identifies the RADIUS server and the AAA server group to which it belongs.</p> <p>When you enter the aaa-server host command, you enter aaa-server host configuration mode.</p>
Step 2	<pre>acl-netmask-convert {auto-detect standard wildcard}</pre> <p>Example: <pre>ciscoasa(config-aaa-server-host)# acl-netmask-convert standard</pre></p>	<p>Specifies how the ASA treats netmasks received in a downloadable ACL from a RADIUS server that is accessed by using the aaa-server host command.</p> <p>The auto-detect keyword specifies that the ASA should attempt to determine the type of netmask expression used. If the ASA detects a wildcard netmask expression, it converts it to a standard netmask expression.</p> <p>The standard keyword specifies that the ASA assumes downloadable ACLs received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed.</p> <p>The wildcard keyword specifies that the ASA assumes downloadable ACLs received from the RADIUS server contain only wildcard netmask expressions and converts them all to standard netmask expressions when the ACLs are downloaded.</p>
Step 3	<pre>radius-common-pw string</pre> <p>Example: <pre>ciscoasa(config-aaa-server-host)# radius-common-pw examplepassword123abc</pre></p>	<p>Specifies a common password to be used for all users who are accessing a RADIUS authorization server through the ASA.</p> <p>The <i>string</i> argument is a case-sensitive, alphanumeric keyword of up to 127 characters to be used as a common password for all authorization transactions with the RADIUS server.</p>
Step 4	<pre>mschapv2-capable</pre> <p>Example: <pre>ciscoasa(config-aaa-server-host)# mschapv2-capable</pre></p>	<p>Enables MS-CHAPv2 authentication requests to the RADIUS server.</p>

	Command	Purpose
Step 5	timeout <i>hh:mm:ss</i> Example: ciscoasa(config-aaa-server-host)# timeout 15	Specifies 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	retry-interval <i>seconds</i> Example: ciscoasa(config-aaa-server-host)# retry-interval 8	Configures the amount of time between retry attempts for a particular AAA server designated in a previous aaa-server host command. The <i>seconds</i> argument specifies the retry interval (1-10 seconds) for the request. This is the time that the ASA waits before retrying a connection request.
Step 7	accounting-mode <i>simultaneous</i> Example: ciscoasa(config-aaa-server-group)# accounting-mode simultaneous	Sends accounting messages to all servers in the group. To restore the default of sending messages only to the active server, enter the accounting-mode single command.
Step 8	authentication-port <i>port</i> Example: ciscoasa(config-aaa-server-host)# authentication-port 1645	Specifies the authentication port as port number 1645, or the server port to be used for authentication of users.
Step 9	accounting-port <i>port</i> Example: ciscoasa(config-aaa-server-host)# accounting-port 1646	Specifies the accounting port as port number 1646, or the server port to be used for accounting for this host.
Step 10	key Example: ciscoasa(config-aaa-host)# key myexamplekey1	Specifies the server secret value 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 value, ask the RADIUS server administrator. The maximum length is 64 characters.

Examples

The following example shows how to add a RADIUS server to an existing RADIUS server group:

```
ciscoasa(config)# aaa-server svrgrp1 protocol radius
ciscoasa(config-aaa-server-group)# aaa-server svrgrp1 host 192.168.3.4
ciscoasa(config-aaa-server-host)# acl-netmask-convert wildcard
ciscoasa(config-aaa-server-host)# radius-common-pw myexamplepasswordabc123
ciscoasa(config-aaa-server-host)# mschapv2-capable
ciscoasa(config-aaa-server-host)# timeout 9
ciscoasa(config-aaa-server-host)# retry-interval 7
ciscoasa(config-aaa-server-host)# accounting-mode simultaneous
ciscoasa(config-aaa-server-host)# authentication-port 1650
ciscoasa(config-aaa-server-host)# authorization-port 1645
ciscoasa(config-aaa-server-host)# key mysecretkeyexampleiceage2
ciscoasa(config-aaa-server-host)# exit
ciscoasa(config)#
```


Monitoring RADIUS Servers

To monitor RADIUS servers, enter one of the following commands:

Command	Purpose
<code>show aaa-server</code>	Shows the configured RADIUS server statistics. To clear the RADIUS server configuration, enter the clear aaa-server statistics command.
<code>show running-config aaa-server</code>	Shows the RADIUS server running configuration. To clear RADIUS server statistics, enter the clear configure aaa-server command.

Additional References

For additional information related to implementing AAA through RADIUS servers, see the “RFCs” section on page 34-20.

RFCs

RFC	Title
2138	<i>Remote Authentication Dial In User Service (RADIUS)</i>
2139	<i>RADIUS Accounting</i>
2548	<i>Microsoft Vendor-specific RADIUS Attributes</i>
2868	<i>RADIUS Attributes for Tunnel Protocol Support</i>

Feature History for RADIUS Servers

Table 34-3 lists each feature change and the platform release in which it was implemented.

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 commands: aaa-server protocol, max-failed-attempts, reactivation-mode, accounting-mode simultaneous, aaa-server host, show aaa-server, show running-config aaa-server, clear aaa-server statistics, authentication-port, accounting-port, retry-interval, acl-netmask-convert, clear configure aaa-server, merge-dacl, radius-common-pw, key.
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](#)
- [Monitoring TACACS+ Servers, page 35-6](#)
- [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.

Table 35-1 Supported TACACS+ Authorization Response Attributes (continued)

Attribute	Description
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_ipp	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 41-36](#).

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 35-5](#)

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 35-5](#).
-

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

	Command	Purpose
Step 1	<p>aaa-server <i>server_tag</i> protocol tacacs+</p> <p>Example: <pre>ciscoasa(config)# aaa-server servergroup1 protocol tacacs+ ciscoasa(config-aaa-server-group)#</pre></p>	<p>Identifies the server group name and the protocol.</p> <p>When you enter the aaa-server protocol command, you enter aaa-server group configuration mode.</p>
Step 2	<p>max-failed-attempts <i>number</i></p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# max-failed-attempts 2</pre></p>	<p>Specifies the maximum number of requests sent to a AAA server in the group before trying the next server. The <i>number</i> argument can range from 1 and 5. The default is 3.</p> <p>If you configured a fallback method using the local database (for management access only), and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (by default), so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see the reactivation-mode command in the next step.</p> <p>If you do not have a fallback method, the ASA continues to retry the servers in the group.</p>
Step 3	<p>reactivation-mode {depletion [deadtime <i>minutes</i>] timed}</p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# reactivation-mode deadtime 20</pre></p>	<p>Specifies the method (reactivation policy) by which failed servers in a group are reactivated.</p> <p>The depletion keyword reactivates failed servers only after all of the servers in the group are inactive.</p> <p>The deadtime <i>minutes</i> keyword-argument pair specifies the amount of time in minutes, between 0 and 1440, that elapses between the disabling of the last server in the group and the subsequent reenabling of all servers. The default is 10 minutes.</p> <p>The timed keyword reactivates failed servers after 30 seconds of down time.</p>

	Command	Purpose
Step 4	accounting-mode simultaneous Example: ciscoasa(config-aaa-server-group)# accounting-mode simultaneous	Sends accounting messages to all servers in the group. To restore the default of sending messages only to the active server, enter the accounting-mode single command.

Examples

The following example shows how to add one TACACS+ group with one primary and one backup server:

```
ciscoasa(config)# aaa-server AuthInbound protocol tacacs+
ciscoasa(config-aaa-server-group)# max-failed-attempts 2
ciscoasa(config-aaa-server-group)# reactivation-mode depletion deadtime 20
ciscoasa(config-aaa-server-group)# exit
ciscoasa(config)# aaa-server AuthInbound (inside) host 10.1.1.1
ciscoasa(config-aaa-server-host)# key TACPlusUauthKey
ciscoasa(config-aaa-server-host)# exit
ciscoasa(config)# aaa-server AuthInbound (inside) host 10.1.1.2
ciscoasa(config-aaa-server-host)# key TACPlusUauthKey2
ciscoasa(config-aaa-server-host)# exit
```

Adding a TACACS+ Server to a Group

To add a TACACS+ server to a group, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	aaa-server <i>server_group</i> [<i>interface_name</i>] host <i>server_ip</i> Example: ciscoasa(config-aaa-server-group)# aaa-server servergroup1 outside host 10.10.1.1	Identifies the TACACS+ server and the server group to which it belongs. When you enter the aaa-server host command, you enter aaa-server host configuration mode.
Step 2	timeout <i>hh:mm:ss</i> Example: ciscoasa(config-aaa-server-host)# timeout 15	Specifies 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.

	Command	Purpose
Step 3	<p>server-port <i>port_number</i></p> <p>Example: ciscoasa(config-aaa-server-host)# server-port 49</p>	Specifies the server port as port number 49, or the TCP port number used by the ASA to communicate with the TACACS+ server.
Step 4	<p>key</p> <p>Example: ciscoasa(config-aaa-host)# key myexamplekey1</p>	Specifies the server secret value used to authenticate the NAS to the TACACS+ server. This value is a case-sensitive, alphanumeric keyword of up to 127 characters, which is the same value as the key on the TACACS+ server. Any characters over 127 are ignored. The key is used between the client and the server to encrypt data between them and must be the same on both the client and server systems. The key cannot contain spaces, but other special characters are allowed.

Monitoring TACACS+ Servers

To monitor TACACS+ servers, enter one of the following commands:

Command	Purpose
show aaa-server	Shows the configured TACACS+ server statistics. To clear the TACACS+ server configuration, enter the clear aaa-server statistics command.
show running-config aaa-server	Shows the TACACS+ server running configuration. To clear TACACS+ server statistics, enter the clear configure aaa-server command.

Feature History for TACACS+ Servers

Table 35-3 lists each feature change and the platform release in which it was implemented.

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 commands: aaa-server protocol, max-failed-attempts, reactivation-mode, accounting-mode simultaneous, aaa-server host, aaa authorization exec authentication-server, server-port, key, clear aaa-server statistics, clear configure aaa-server, show aaa-server, show running-config aaa-server, username, service-type, timeout.



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](#)
- [Monitoring LDAP Servers, page 36-11](#)
- [Feature History for LDAP Servers, page 36-12](#)

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**

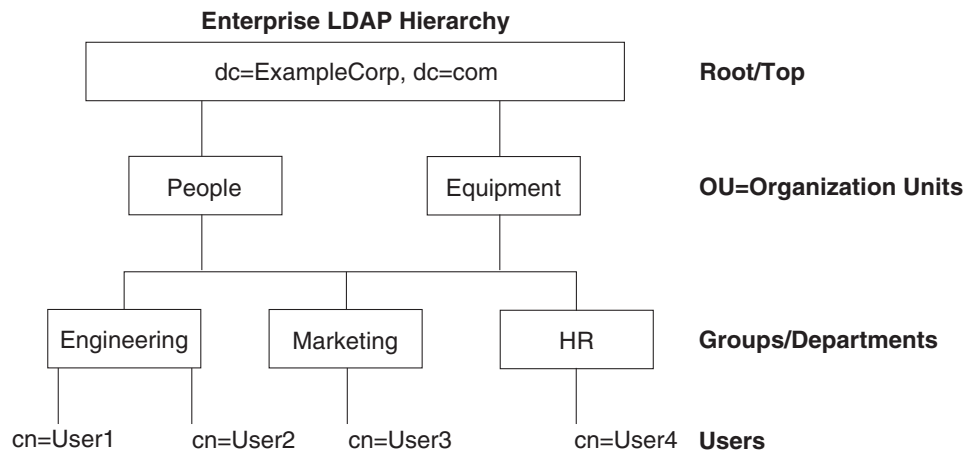
For more information about the LDAP protocol, see RFCs 1777, 2251, and 2849.

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](#)[Configuring Authorization with LDAP for VPN, page 36-10](#)

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 | (Optional) Configure authorization from an LDAP server that is separate and distinct from the authentication mechanism. See the “ Configuring Authorization with LDAP for VPN ” section on page 36-10. |
| 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

	Command	Purpose
Step 1	ldap attribute-map <i>map-name</i> Example: ciscoasa(config)# ldap attribute-map att_map_1	Creates an unpopulated LDAP attribute map table.
Step 2	map-name <i>user-attribute-name</i> <i>Cisco-attribute-name</i> Example: ciscoasa(config-ldap-attribute-map)# map-name department IETF-Radius-Class	Maps the user-defined attribute name department to the Cisco attribute.
Step 3	map-value <i>user-attribute-name</i> <i>Cisco-attribute-name</i> Example: ciscoasa(config-ldap-attribute-map)# map-value department Engineering group1	Maps the user-defined map value department to the user-defined attribute value and the Cisco attribute value.
Step 4	aaa-server <i>server_group</i> [<i>interface_name</i>] host <i>server_ip</i> Example: ciscoasa(config)# aaa-server ldap_dir_1 host 10.1.1.4	Identifies the server and the AAA server group to which it belongs.
Step 5	ldap-attribute-map <i>map-name</i> Example: ciscoasa(config-aaa-server-host)# ldap-attribute-map att_map_1	Binds the attribute map to the LDAP server.

Examples

The following example shows how to limit management sessions to the ASA based on an LDAP attribute called `accessType`. The `accessType` attribute may have one of these values:

- VPN
- admin
- helpdesk

The following example shows how each value is mapped to one of the valid IETF-RADIUS-Service-Type attributes that the ASA supports: `remote-access` (Service-Type 5) Outbound, `admin` (Service-Type 6) Administrative, and `nas-prompt` (Service-Type 7) NAS Prompt.

```
ciscoasa(config)# ldap attribute-map MGMT
ciscoasa(config-ldap-attribute-map)# map-name accessType IETF-RADIUS-Service-Type
ciscoasa(config-ldap-attribute-map)# map-value accessType VPN 5
ciscoasa(config-ldap-attribute-map)# map-value accessType admin 6
ciscoasa(config-ldap-attribute-map)# map-value accessType helpdesk 7

ciscoasa(config-ldap-attribute-map)# aaa-server LDAP protocol ldap
ciscoasa(config-aaa-server-group)# aaa-server LDAP (inside) host 10.1.254.91
ciscoasa(config-aaa-server-host)# ldap-base-dn CN=Users,DC=cisco,DC=local
ciscoasa(config-aaa-server-host)# ldap-scope subtree
ciscoasa(config-aaa-server-host)# ldap-login-password test
ciscoasa(config-aaa-server-host)# ldap-login-dn
CN=Administrator,CN=Users,DC=cisco,DC=local
ciscoasa(config-aaa-server-host)# server-type auto-detect
ciscoasa(config-aaa-server-host)# ldap-attribute-map MGMT
```

The following example shows how to display the complete list of Cisco LDAP attribute names:

```
ciscoasa(config)# ldap attribute-map att_map_1
ciscoasa(config-ldap-attribute-map)# map-name att_map_1?

ldap mode commands/options:
cisco-attribute-names:
  Access-Hours
  Allow-Network-Extension-Mode
  Auth-Service-Type
  Authenticated-User-Idle-Timeout
  Authorization-Required
  Authorization-Type
  :
  :
  X509-Cert-Data
ciscoasa(config-ldap-attribute-map)#
```

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.

	Command	Purpose
Step 1	<p>aaa-server <i>server_tag</i> protocol ldap</p> <p>Example: <pre>ciscoasa(config)# aaa-server servergroup1 protocol ldap ciscoasa(config-aaa-server-group)#</pre></p>	<p>Identifies the server group name and the protocol. When you enter the aaa-server protocol command, you enter aaa-server group configuration mode.</p>
Step 2	<p>max-failed-attempts <i>number</i></p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# max-failed-attempts 2</pre></p>	<p>Specifies the maximum number of requests sent to an LDAP server in the group before trying the next server. The <i>number</i> argument can range from 1 and 5. The default is 3.</p> <p>If you configured a fallback method using the local database (for management access only) to configure the fallback mechanism), and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (by default), so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see the reactivation-mode command in the next step.</p> <p>If you do not have a fallback method, the ASA continues to retry the servers in the group.</p>

	Command	Purpose
Step 3	<pre>reactivation-mode {depletion [deadtime minutes] timed} Example: ciscoasa(config-aaa-server-group)# reactivation-mode deadtime 20</pre>	<p>Specifies the method (reactivation policy) by which failed servers in a group are reactivated.</p> <p>The depletion keyword reactivates failed servers only after all of the servers in the group are inactive.</p> <p>The deadtime minutes keyword-argument pair specifies the amount of time in minutes, between 0 and 1440, that elapses between the disabling of the last server in the group and the subsequent reenabling of all servers. The default is 10 minutes.</p> <p>The timed keyword reactivates failed servers after 30 seconds of down time.</p>
Step 4	<pre>aaa-server server_group [interface_name] host server_ip Example: ciscoasa(config)# aaa-server servergroup1 outside host 10.10.1.1 Move to new procedure for adding a server to a groupyp</pre>	<p>Identifies the LDAP server and AAA server group to which it belongs.</p> <p>When you enter the aaa-server host command, you enter aaa-server host configuration mode. As needed, use host configuration mode commands to further configure the AAA server.</p> <p>Table 36-2 lists the available commands for LDAP servers, and whether or not a new LDAP server definition has a default value for that command. If no default value is provided (indicated by “—”), use the command to specify the value.</p>

Table 36-2 Host Mode Commands and Defaults

Command	Default Value	Description
ldap-attribute-map	—	Separate steps in procedure under host command
ldap-base-dn	—	
ldap-login-dn	—	
ldap-login-password	—	
ldap-naming-attribute	—	
ldap-over-ssl	636	If not set, the ASA uses sAMAccountName for LDAP requests. Whether using SASL or plain text, you can secure communications between the ASA and the LDAP server with SSL. If you do not configure SASL, we strongly recommend that you secure LDAP communications with SSL.
ldap-scope	—	
sasl-mechanism	—	
server-port	389	
server-type	autodiscovery	If autodetection fails to determine the LDAP server type, and you know the server is either a Microsoft, Sun or generic LDAP server, you can manually configure the server type.
timeout	10 seconds	

Examples

The following example shows how to configure an LDAP server group named `watchdogs` and add an LDAP server to the group. Because the example does not define a retry interval or the port that the LDAP server listens to, the ASA uses the default values for these two server-specific parameters.

```
ciscoasa(config)# aaa-server watchdogs protocol ldap
ciscoasa(config-aaa-server-group)# aaa-server watchdogs host 192.168.3.4
ciscoasa(config-aaa-server-host)# exit
ciscoasa(config)#
```

Configuring Authorization with LDAP for VPN

When user LDAP authentication for VPN access has succeeded, the ASA queries the LDAP server, which returns LDAP attributes. These attributes generally include authorization data that applies to the VPN session. Using LDAP in this way accomplishes authentication and authorization in a single step.

There may be cases, however, where you require authorization from an LDAP directory server that is separate and distinct from the authentication mechanism. For example, if you use an SDI or certificate server for authentication, no authorization information is returned. For user authorizations in this case, you can query an LDAP directory after successful authentication, accomplishing authentication and authorization in two steps.

To set up VPN user authorization using LDAP, perform the following steps.

Detailed Steps

	Command	Purpose
Step 1	tunnel-group <i>groupname</i> Example: ciscoasa(config)# tunnel-group remotegrp	Creates an IPsec remote access tunnel group named remotegrp.
Step 2	tunnel-group <i>groupname</i> general-attributes Example: ciscoasa(config)# tunnel-group remotegrp general-attributes	Associates the server group and the tunnel group.
Step 3	authorization-server-group <i>group-tag</i> Example: ciscoasa(config-general)# authorization-server-group ldap_dir_1	Assigns a new tunnel group to a previously created AAA server group for authorization.

Examples

While there are other authorization-related commands and options available for specific requirements, the following example shows commands for enabling user authorization with LDAP. The example then creates an IPsec remote access tunnel group named `remote-1`, and assigns that new tunnel group to the previously created `ldap_dir_1` AAA server group for authorization:

```
ciscoasa(config)# tunnel-group remote-1 type ipsec-ra
```

```
ciscoasa(config)# tunnel-group remote-1 general-attributes
ciscoasa(config-general)# authorization-server-group ldap_dir_1
ciscoasa(config-general)#
```

After you complete this configuration work, you can then configure additional LDAP authorization parameters such as a directory password, a starting point for searching a directory, and the scope of a directory search by entering the following commands:

```
ciscoasa(config)# aaa-server ldap_dir_1 protocol ldap
ciscoasa(config-aaa-server-group)# aaa-server ldap_dir_1 host 10.1.1.4
ciscoasa(config-aaa-server-host)# ldap-login-dn obscurepassword
ciscoasa(config-aaa-server-host)# ldap-base-dn starthere
ciscoasa(config-aaa-server-host)# ldap-scope subtree
ciscoasa(config-aaa-server-host)#
```

Monitoring LDAP Servers

To monitor LDAP servers, enter one of the following commands:

Command	Purpose
show aaa-server	Shows the configured AAA server statistics. To clear the AAA server configuration, enter the clear aaa-server statistics command.
show running-config aaa-server	Shows the AAA server running configuration. To clear AAA server statistics, enter the clear configure aaa-server command.

Feature History for LDAP Servers

Table 36-3 lists each feature change and the platform release in which it was implemented.

Table 36-3 Feature History for AAA Servers

Feature Name	Platform Releases	Feature Information
LDAP Servers for AAA	7.0(1)	<p>LDAP Servers describe support for AAA and how to configure LDAP servers.</p> <p>We introduced the following commands:</p> <p>username, aaa authorization exec authentication-server, aaa authentication console LOCAL, aaa authorization exec LOCAL, service-type, ldap attribute-map, aaa-server protocol, aaa authentication {telnet ssh serial} console LOCAL, aaa authentication http console LOCAL, aaa authentication enable console LOCAL, max-failed-attempts, reactivation-mode, accounting-mode simultaneous, aaa-server host, authorization-server-group, tunnel-group, tunnel-group general-attributes, map-name, map-value, ldap-attribute-map.</p>



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](#)
- [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 37-4](#)

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 37-4 . |
-

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

	Command	Purpose
Step 1	<p>aaa-server <i>server_tag</i> protocol nt</p> <p>Example: <pre>ciscoasa(config)# aaa-server servergroup1 protocol nt ciscoasa(config-aaa-server-group)#</pre></p>	<p>Identifies the server group name and the protocol.</p> <p>When you enter the aaa-server protocol command, you enter aaa-server group configuration mode.</p>
Step 2	<p>max-failed-attempts <i>number</i></p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# max-failed-attempts 2</pre></p>	<p>Specifies the maximum number of requests sent to a Windows NT server in the group before trying the next server. The <i>number</i> argument can range from 1 and 5. The default is 3.</p> <p>If you configured a fallback method using the local database (for management access only), and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (by default), so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see the reactivation-mode command in the next step.</p> <p>If you do not have a fallback method, the ASA continues to retry the servers in the group.</p>
Step 3	<p>reactivation-mode {depletion [<i>deadtime minutes</i>] timed}</p> <p>Example: <pre>ciscoasa(config-aaa-server-group)# reactivation-mode deadtime 20</pre></p>	<p>Specifies the method (reactivation policy) by which failed servers in a group are reactivated.</p> <p>The depletion keyword reactivates failed servers only after all of the servers in the group are inactive.</p> <p>The deadtime minutes keyword-argument pair specifies the amount of time in minutes, between 0 and 1440, that elapses between the disabling of the last server in the group and the subsequent reenabling of all servers. The default is 10 minutes.</p> <p>The timed keyword reactivates failed servers after 30 seconds of down time.</p>

Examples

The following example shows how to add a Windows NT domain server group:

```
ciscoasa(config)# aaa-server NTAuth protocol nt
ciscoasa(config-aaa-server-group)# max-failed-attempts 2
ciscoasa(config-aaa-server-group)# reactivation-mode depletion deadline 20
ciscoasa(config-aaa-server-group)# exit
ciscoasa(config)# aaa-server NTAuth (inside) host 10.1.1.4
ciscoasa(config-aaa-server-host)# nt-auth-domain-controller primary1
ciscoasa(config-aaa-server-host)# exit
```

Adding a Windows NT Server to a Group

To add a Windows NT server to a group, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>aaa-server <i>server_group</i> [<i>interface_name</i>] host <i>server_ip</i></pre> <p>Example: <pre>ciscoasa(config-aaa-server-group)# aaa-server servergroup1 outside host 10.10.1.1</pre></p>	<p>Identifies the Windows NT server and the server group to which it belongs.</p> <p>When you enter the aaa-server host command, you enter aaa-server host configuration mode.</p>
Step 2	<pre>timeout <i>hh:mm:ss</i></pre> <p>Example: <pre>ciscoasa(config-aaa-server-host)# timeout 15</pre></p>	<p>Specifies the length of time, in hours, minutes, and seconds, that the ASA waits for a response from the primary server before sending the request to the backup server.</p>
Step 3	<pre>server-port <i>port_number</i></pre> <p>Example: <pre>ciscoasa(config-aaa-server-host)# server-port 139</pre></p>	<p>Specifies the server port as port number 139, or the TCP port number used by the ASA to communicate with the Windows NT server.</p>
Step 4	<pre>nt-auth-domain-controller <i>string</i></pre> <p>Example: <pre>ciscoasa(config-aaa-server)# nt-auth-domain controller primary1</pre></p>	<p>Specifies the name for the Windows NT authentication domain controller.</p> <p>The <i>string</i> argument represents 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.</p>

Examples

The following example shows how to add a Windows NT domain server to the NTAAuth server group:

```
ciscoasa(config)# aaa-server NTAuth (inside) host 10.1.1.4
ciscoasa(config-aaa-server-host)# timeout 15
```

```
ciscoasa(config-aaa-server-host)# server-port 139
ciscoasa(config-aaa-server-host)# nt-auth-domain-controller primary1
ciscoasa(config-aaa-server-host)# exit
```

Monitoring Windows NT Servers

To monitor Windows NT servers, enter one of the following commands:

Command	Purpose
show aaa-server	Shows the configured Windows NT server statistics. To clear the Windows NT server statistics, enter the clear aaa-server statistics command.
show running-config aaa-server	Shows the Windows NT server running configuration. To clear Windows NT server configuration, enter the clear configure aaa-server command.

Feature History for Windows NT Servers

[Table 37-1](#) lists each feature change and the platform release in which it was implemented.

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 commands: aaa-server protocol, max-failed-attempts, clear configure aaa-server, clear aaa-server statistics, reactivation-mode, aaa-server host, server-port, timeout, nt-auth-domain-controller, show aaa-server, show running-config aaa-server.



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 38-23](#)
- [Feature History for the Identity Firewall, page 38-25](#)

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

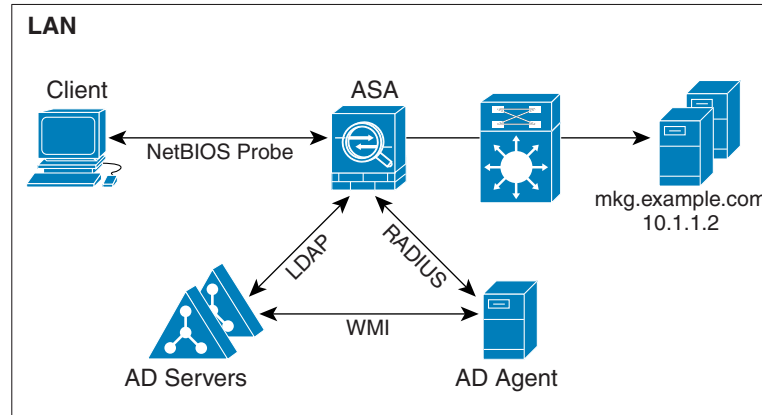
The AD Agent runs on a Windows server. Supported Windows servers include Windows 2003, Windows 2008, and Windows 2008 R2.



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	<p>On the ASA: Administrators configure local user groups and Identity Firewall policies.</p>	4	<p>Client <-> ASA: The client logs into the network through Microsoft Active Directory. The AD Server authenticates users and generates user login security logs.</p> <p>Alternatively, the client can log into the network through a cut-through proxy or VPN.</p>
2	<p>ASA <-> AD Server: The ASA sends an LDAP query for the Active Directory groups configured on the AD Server.</p> <p>The ASA consolidates local and Active Directory groups and applies access rules and Modular Policy Framework security policies based on user identity.</p>	5	<p>ASA <-> Client: Based on the policies configured on the ASA, it grants or denies access to the client.</p> <p>If configured, the ASA probes the NetBIOS of the client to pass inactive and no-response users.</p>
3	<p>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.</p> <p>The ASA forwards the new mapped entries that have been learned from web authentication and VPN sessions to the AD Agent.</p>	6	<p>AD Agent <-> AD Server: The AD Agent maintains a cache of user ID and IP address mapped entries, and notifies the ASA of changes.</p> <p>The AD Agent sends logs to a syslog server.</p>

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.

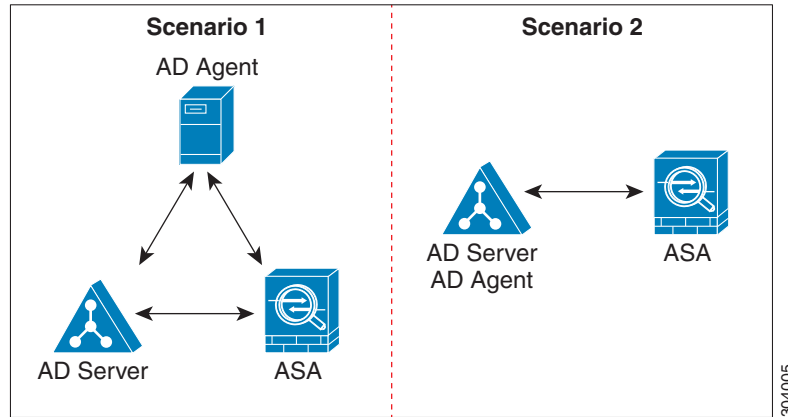
Figure 38-2 Deployment Scenario without Redundancy

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.

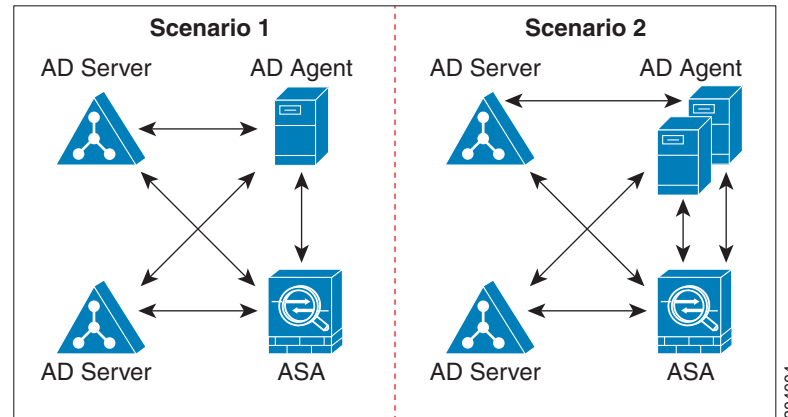
Figure 38-3 Deployment Scenario with Redundant Components

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.

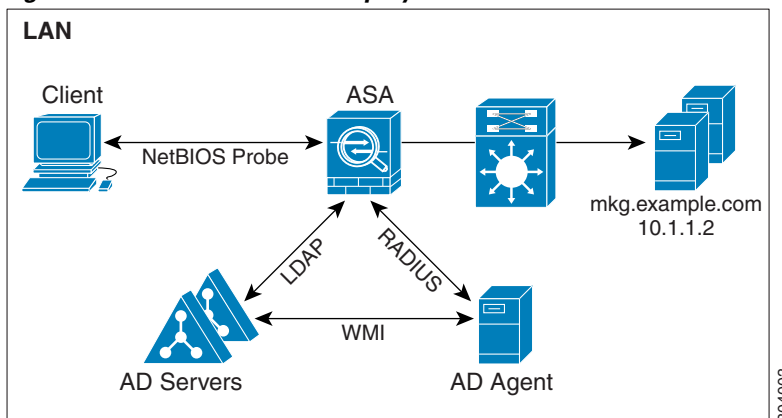
Figure 38-4 LAN-based Deployment

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.

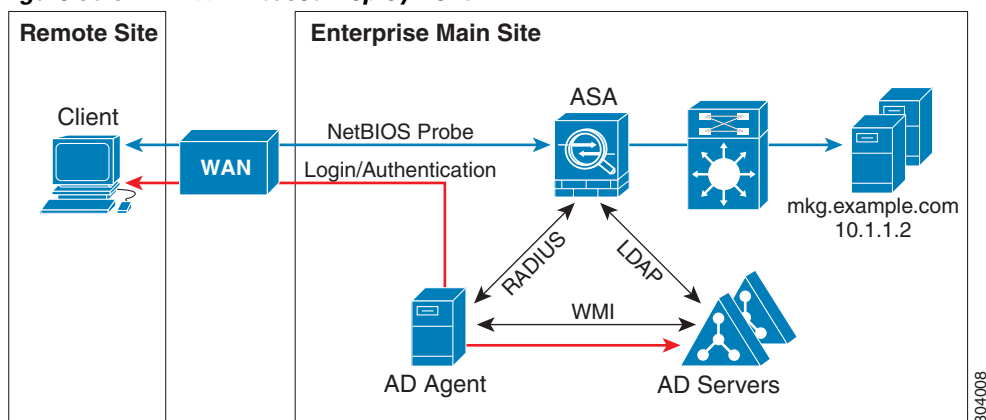
Figure 38-5 WAN-based Deployment

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.

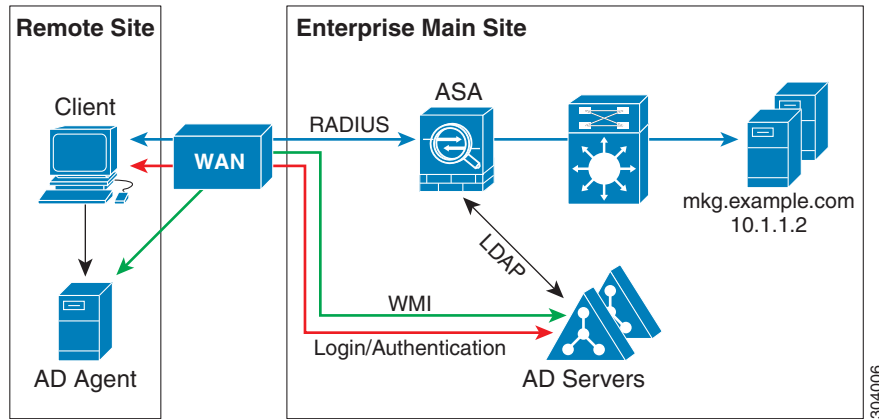
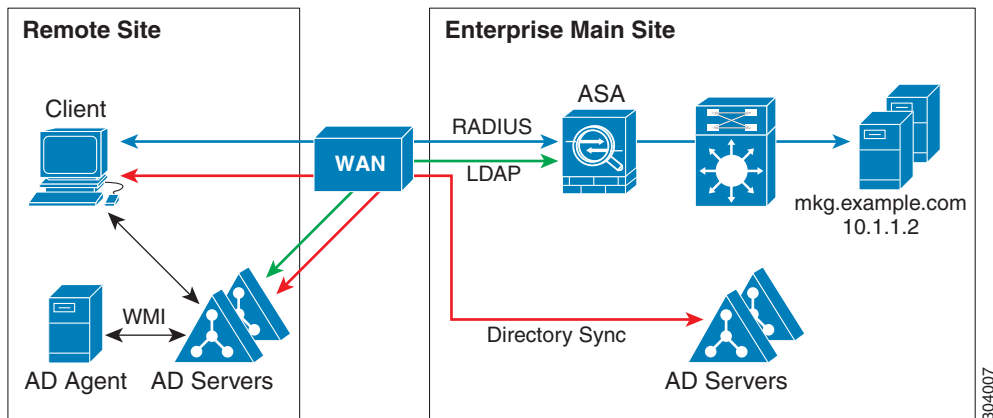
Figure 38-6 WAN-based Deployment with Remote AD Agent

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: V:*?"<>|. For naming conventions, see <http://support.microsoft.com/kb/909264>.
- For usernames, the following characters are not valid: V[!];=,*?"<>|@.
- For user group names, the following characters are not valid: V[!];=,*?"<>|.

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 topics:

- [Task Flow for Configuring the Identity Firewall](#), page 38-10
- [Configuring the Active Directory Domain](#), page 38-11
- [Configuring Active Directory Agents](#), page 38-13
- [Configuring Identity Options](#), page 38-14
- [Configuring Identity-Based Security Policy](#), page 38-19
- [Collecting User Statistics](#), page 38-20

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.
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 Agents](#)” 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 38-14.

- 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 38-19.

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:

	Command	Purpose
Step 1	<code>aaa-server server-tag protocol ldap</code> Example: hostname(config)# aaa-server adserver protocol ldap	Creates the AAA server group and configures AAA server parameters for the Active Directory server.
Step 2	<code>aaa-server server-tag [(interface-name)] host {server-ip name} [key] [timeout seconds]</code> Example: hostname(config-aaa-server-group)# aaa-server adserver (mgmt) host 172.168.224.6	For the Active Directory server, configures the AAA server as part of a AAA server group and the AAA server parameters that are host-specific.
Step 3	<code>ldap-base-dn string</code> Example: hostname(config-aaa-server-host)# ldap-base-dn DC=SAMPLE,DC=com	Specifies the location in the LDAP hierarchy where the server should begin searching when it receives an authorization request. Specifying the ldap-base-dn command is optional. If you do not specify this command, the ASA retrieves the defaultNamingContext from the Active Directory and uses it as the base DN.
Step 4	<code>ldap-scope subtree</code> Example: hostname(config-aaa-server-host)# ldap-scope subtree	Specifies the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request.

	Command	Purpose
Step 5	<p>ldap-login-password <i>string</i></p> <p>Example: hostname(config-aaa-server-host)# ldap-login-password obscurepassword</p>	Specifies the login password for the LDAP server.
Step 6	<p>ldap-login-dn <i>string</i></p> <p>Example: hostname(config-aaa-server-host)# ldap-login-dn SAMPLE\user1</p>	<p>Specifies the name of the directory object that the system should bind this as. The ASA identifies itself for authenticated binding by attaching a Login DN field to the user authentication request. The Login DN field describes the authentication characteristics of the ASA.</p> <p>The <i>string</i> argument is a case-sensitive string of up to 128 characters that specifies the name of the directory object in the LDAP hierarchy. Spaces are not permitted in the string, but other special characters are allowed.</p> <p>You can specify the traditional or simplified format.</p> <p>The typical ldap-login-dn command format includes: CN=username,OU=Employees,OU=Sample Users,DC=sample,DC=com.</p>
Step 7	<p>server-type <i>microsoft</i></p> <p>Example: hostname(config-aaa-server-host)# server-type microsoft</p>	Configures the LDAP server model for the Microsoft Active Directory server.
Step 8	<p>ldap-group-base-dn <i>string</i></p> <p>Example: hostname(config-aaa-server-host)# ldap-group-base-dn OU=Sample Groups,DC=SAMPLE,DC=com</p>	<p>Specifies location of the Active Directory groups configuration in the Active Directory domain controller. If not specified, the value in the ldap-group-base-dn command is used.</p> <p>Specifying the ldap-group-base-dn command is optional.</p>
Step 9	<p>ldap-over-ssl <i>enable</i></p> <p>Example: hostname(config-aaa-server-host)# ldap-over-ssl enable</p>	<p>Allows the ASA to access the Active Directory domain controller over SSL. To support LDAP over SSL, Active Directory server needs to be configured to have this support.</p> <p>By default, the Active Directory does not have SSL configured. If SSL is not configured on the Active Directory, you do not need to configure it on the ASA for the Identity Firewall.</p>

	Command	Purpose
Step 10	server-port <i>port-number</i> Example: hostname(config-aaa-server-host)# server-port 389 hostname(config-aaa-server-host)# server-port 636	By default, if the ldap-over-ssl command is not enabled, the default server port is 389; if the ldap-over-ssl command is enabled, the default server port is 636.
Step 11	group-search-timeout <i>seconds</i> Example: hostname(config-aaa-server-host)# group-search-timeout 300	Sets the amount of time before LDAP queries time out.

Configuring Active Directory Agents

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.

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:

	Command	Purpose
Step 1	aaa-server <i>server-tag</i> protocol radius Example: hostname(config)# aaa-server adagent protocol radius	Creates the AAA server group and configures AAA server parameters for the AD Agent.
Step 2	ad-agent-mode Example: hostname(config)# ad-agent-mode	Enables the AD Agent mode.
Step 3	aaa-server <i>server-tag</i> [(<i>interface-name</i>)] host { <i>server-ip</i> <i>name</i> } [<i>key</i>] [timeout <i>seconds</i>] Example: hostname(config-aaa-server-group)# aaa-server adagent (inside) host 192.168.1.101	For the AD Agent, configures the AAA server as part of a AAA server group and the AAA server parameters that are host-specific.

	Command	Purpose
Step 4	<p>key <i>key</i></p> <p>Example: hostname(config-aaa-server-host)# key mysecret</p>	Specifies the server secret value used to authenticate the ASA to the AD Agent server.
Step 5	<p>user-identity ad-agent aaa-server <i>aaa_server_group_tag</i></p> <p>Example: hostname(config-aaa-server-hostkey)# user-identity ad-agent aaa-server adagent</p>	<p>Defines the server group of the AD Agent.</p> <p>The first server defined in the <i>aaa_server_group_tag</i> argument is the primary AD Agent and the second server defined is the secondary AD Agent.</p> <p>The Identity Firewall supports defining only two AD Agent hosts.</p> <p>When the ASA detects that the primary AD Agent is down and a secondary agent is specified, it switches to the secondary AD Agent. The AAA server for the AD agent uses RADIUS as the communication protocol, and should specify a key attribute for the shared secret between the ASA and AD Agent.</p>
Step 6	<p>test aaa-server ad-agent</p> <p>Example: hostname(config-aaa-server-host)# test aaa-server ad-agent</p>	Tests the communication between the ASA and the AD Agent server.

What to Do Next

Configure access rules for the Identity Firewall. See the [“Configuring Identity-Based Security Policy” section on page 38-19](#).

Configuring Identity Options

Perform this procedure 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:

	Command	Purpose
Step 1	<p>user-identity enable</p> <p>Example: hostname(config)# user-identity enable</p>	Enables the Identity Firewall feature.
Step 2	<p>user-identity default-domain domain_NetBIOS_name</p> <p>Example: hostname(config)# user-identity default-domain SAMPLE</p>	<p>Specifies the default domain for the Identity Firewall.</p> <p>For the <i>domain_NetBIOS_name</i> argument, enter a name of up to 32 characters that consists of [a-z], [A-Z], [0-9], [!@#%&()-_+=[]{};,.] except '.' and '' at the first character. If the domain name includes a space, enclose the entire name in quotation marks. The domain name is not case sensitive.</p> <p>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. For multiple context modes, you can set a default domain name for each context, as well as within the system execution space.</p> <p>Note The default domain name that you specify 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 identity-IP address mapped entries with the domain name that you enter when configuring the ASA. To view the NetBIOS domain name, open the Active Directory user event security log in any text editor.</p> <p>The Identity Firewall uses the LOCAL domain for all locally defined user groups or locally defined users. Users logging in through a web portal (cut-through proxy) are designated as belonging to the Active Directory domain with which they authenticated. Users logging in through a VPN are designated as belonging to the LOCAL domain unless the VPN is authenticated by LDAP with the Active Directory. In this case, the Identity Firewall can associate the users with their Active Directory domain.</p>

	Command	Purpose
Step 3	<pre>user-identity domain domain_nickname aaa-server aaa_server_group_tag</pre> <p>Example:</p> <pre>hostname(config)# user-identity domain SAMPLE aaa-server ds</pre>	<p>Associates the LDAP parameters defined for the AAA server for importing user group queries with the domain name.</p> <p>For the <i>domain_nickname</i> argument, enter a name of 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.</p>
Step 4	<pre>user-identity logout-probe netbios local-system probe-time minutes minutes retry-interval seconds seconds retry-count times [user-not-needed match-any exact-match]</pre> <p>Example:</p> <pre>hostname(config)# user-identity logout-probe netbios local-system probe-time minutes 10 retry-interval seconds 10 retry-count 2 user-not-needed</pre>	<p>Enables NetBIOS probing. Enabling this option configures how often the ASA probes the user client IP address to determine whether the client is still active. By default, NetBIOS probing is disabled.</p> <p>To minimize the NetBIOS packets, the ASA only sends a NetBIOS probe to a client when the user has been idle for more than the specified number of minutes.</p> <ul style="list-style-type: none"> • 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. • User-not-needed—As long as the ASA received a NetBIOS response from the client, the user identity is considered valid. <p>The Identity Firewall only performs NetBIOS probing for those users identities that are in the active state and exist in at least one security policy. The ASA does not perform NetBIOS probing for clients where the users logged in through cut-through proxy or by using a VPN.</p>
Step 5	<pre>user-identity inactive-user-timer minutes minutes</pre> <p>Example:</p> <pre>hostname(config)# user-identity inactive-user-timer minutes 120</pre>	<p>Specifies the amount of time before a user is considered idle, meaning the ASA has not received traffic from the user's IP address for the specified amount of time.</p> <p>When the timer expires, the user's IP address is marked as inactive and removed from the local cached user identity-IP address mapping database, and the ASA no longer notifies the AD Agent about that IP address. Existing traffic is still allowed to pass. When this command is specified, the ASA runs an inactive timer even when the NetBIOS Logout Probe is configured.</p> <p>By default, the idle timeout is set to 60 minutes.</p> <p>Note The Idle Timeout option does not apply to VPN or cut-through proxy users.</p>

	Command	Purpose
Step 6	<pre>user-identity poll-import-user-group-timer hours hours</pre> <p>Example:</p> <pre>hostname(config)# user-identity poll-import-user-group-timer hours 1</pre>	<p>Specifies the amount of time before the ASA queries the Active Directory server for user group information.</p> <p>If a user is added to or deleted from an Active Directory group, the ASA received the updated user group after the import group timer ran.</p> <p>By default, the poll-import-user-group-timer hours value is 8 hours.</p> <p>To immediately update user group information, enter the user-identity update import-user command.</p>
Step 7	<pre>user-identity action netbios-response-fail remove-user-ip</pre> <p>Example:</p> <pre>hostname(config)# user-identity action netbios-response-fail remove-user-ip</pre>	<p>Specifies the action when a client does not respond to a NetBIOS probe. For example, the network connection might be blocked to that client or the client is not active.</p> <p>When the user-identity action remove-user-ip command is configured, the ASA removed the user identity-IP address mapping for that client.</p> <p>By default, this command is disabled.</p>
Step 8	<pre>user-identity action domain-controller-down domain_nickname disable-user-identity-rule</pre> <p>Example:</p> <pre>hostname(config)# user-identity action domain-controller-down SAMPLE disable-user-identity-rule</pre>	<p>Specifies the action when the domain is down, because the Active Directory domain controller is not responding.</p> <p>When the domain is down and the disable-user-identity-rule keyword is configured, the ASA disables the user identity-IP address mapping for that domain. Additionally, the status of all user IP addresses in that domain are marked as disabled in the output displayed by the show user-identity user command.</p> <p>By default, this command is disabled.</p>
Step 9	<pre>user-identity user-not-found enable</pre> <p>Example:</p> <pre>hostname(config)# user-identity user-not-found enable</pre>	<p>Enables user-not-found tracking. Only the last 1024 IP addresses are tracked.</p> <p>By default, this command is disabled.</p>
Step 10	<pre>user-identity action ad-agent-down disable-user-identity-rule</pre> <p>Example:</p> <pre>hostname(config)# user-identity action ad-agent-down disable-user-identity-rule</pre>	<p>Specifies the action when the AD Agent is not responding.</p> <p>When the AD Agent is down and the user-identity action ad-agent-down command is configured, 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 output displayed by the show user-identity user command.</p> <p>By default, this command is disabled.</p>

	Command	Purpose
Step 11	<pre>user-identity action mac-address-mismatch remove-user-ip</pre> <p>Example:</p> <pre>hostname(config)# user-identity action mac-address-mismatch remove-user-ip</pre>	<p>Specifies the action when a user's MAC address is found to be inconsistent with the ASA IP address currently mapped to that MAC address.</p> <p>When the user-identity action mac-address-mismatch command is configured, the ASA removes the user identity-IP address mapping for that client.</p> <p>By default, the ASA uses the remove-user-ip keyword when this command is specified.</p>
Step 12	<pre>user-identity ad-agent active-user-database {on-demand full-download}</pre> <p>Example:</p> <pre>hostname(config)# user-identity ad-agent active-user-database full-download</pre>	<p>Defines how the ASA retrieves the user identity-IP address mapping information from the AD Agent:</p> <ul style="list-style-type: none"> • 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 information when users log in and log out. • 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. <p>By default, the ASA 5505 uses the on-demand option. The other ASA platforms use the full-download option.</p> <p>Full downloads are event driven, meaning that when there are subsequent requests to download the database, just the updates to the user identity-IP address mapping database are sent.</p> <p>When the ASA registers a change request with the AD Agent, the AD Agent sends a new event to the ASA.</p>

	Command	Purpose
Step 13	<pre>user-identity ad-agent hello-timer seconds <i>seconds</i> retry-times <i>number</i></pre> <p>Example: <pre>hostname(config)# user-identity ad-agent hello-timer seconds 20 retry-times 3</pre></p>	<p>Defines the hello timer between the ASA and the AD Agent.</p> <p>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.</p> <p>By default, the hello timer is set to 30 seconds and 5 retries.</p>
Step 14	<pre>user-identity ad-agent aaa-server aaa_server_group_tag</pre> <p>Example: <pre>hostname(config)# user-identity ad-agent aaa-server adagent</pre></p>	<p>Defines the server group of the AD Agent.</p> <p>For the <i>aaa_server_group_tag</i> argument, enter the value defined by the aaa-server command.</p>

What to Do Next

Configure the Active Directory domain and server groups. See the [“Configuring the Active Directory Domain” section on page 38-11](#).

Configure AD Agents. See the [“Configuring Active Directory Agents” section on page 38-13](#).

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.

- To configure an extended ACL, see [Chapter 19, “Adding an Extended Access Control List.”](#)
- To configure local user groups, which can be used in the ACL, see the [“Configuring Local User Groups” section on page 17-11](#).

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.

Collecting User Statistics

To activate the collection of user statistics by the Modular Policy Framework and match lookup actions for the Identify Firewall, enter the following command:

Command	Purpose
<pre>user-statistics [accounting scanning]</pre> <p>Example:</p> <pre>ciscoasa(config)# class-map c-identity-example-1 ciscoasa(config-cmap)# match access-list identity-example-1 ciscoasa(config-cmap)# exit ciscoasa(config)# policy-map p-identity-example-1 ciscoasa(config-pmap)# class c-identity-example-1 ciscoasa(config-pmap)# user-statistics accounting ciscoasa(config-pmap)# exit ciscoasa(config)# service-policy p-identity-example-1 interface outside</pre>	<p>Activates the collection of user statistics by the Modular Policy Framework and matches lookup actions for the Identify Firewall.</p> <p>The accounting keyword specifies that the ASA collect the sent packet count, sent drop count, and received packet count. The scanning keyword specifies that the ASA collect only the sent drop count.</p> <p>When you configure a policy map to collect user statistics, the ASA collects detailed statistics for selected users. When you specify the user-statistics command without the accounting or scanning keywords, the ASA collects both accounting and scanning statistics.</p>

Configuration Examples

This section includes the following topics:

- [AAA Rule and Access Rule Example 1, page 38-21](#)
- [AAA Rule and Access Rule Example 2, page 38-21](#)
- [VPN Filter Example, page 38-22](#)

AAA Rule and Access Rule Example 1

This example shows a typical cut-through proxy configuration to allow a user to log in through the ASA. In this example, the following conditions apply:

- The ASA IP address is 172.1.1.118.
- The Active Directory domain controller has the IP address 71.1.2.93.
- The end-user client has the IP address 172.1.1.118 and uses HTTPS to log in through a web portal.
- The user is authenticated by the Active Directory domain controller via LDAP.
- The ASA uses the inside interface to connect to the Active Directory domain controller on the corporate network.

```
hostname(config)# access-list AUTH extended permit tcp any 172.1.1.118 255.255.255.255 eq http
hostname(config)# access-list AUTH extended permit tcp any 172.1.1.118 255.255.255.255 eq https
hostname(config)# aaa-server LDAP protocol ldap
hostname(config-aaa-server-group)# aaa-server LDAP (inside) host 171.1.2.93
hostname(config-aaa-server-host)# ldap-base-dn DC=cisco,DC=com
hostname(config-aaa-server-host)# ldap-group-base-dn DC=cisco,DC=com
hostname(config-aaa-server-host)# ldap-scope subtree
hostname(config-aaa-server-host)# ldap-login-dn cn=kao,OU=Employees,OU=Cisco Users,DC=cisco,DC=com
hostname(config-aaa-server-host)# ldap-login-password *****
hostname(config-aaa-server-host)# ldap-over-ssl enable
hostname(config-aaa-server-host)# server-type microsoft
hostname(config-aaa-server-host)# aaa authentication match AUTH inside LDAP
hostname(config)#
hostname(config)# http server enable
hostname(config)# http 0.0.0.0 0.0.0.0 inside
hostname(config)#
hostname(config)# auth-prompt prompt Enter Your Authentication
hostname(config)# auth-prompt accept You are Good
hostname(config)# auth-prompt reject Goodbye
```

AAA Rule and Access Rule Example 2

In this example, the following guidelines apply:

- In **access list** commands, permit user NONE rules should be written before entering the **access-list 100 ex deny any any** command to allow unauthenticated incoming users to trigger AAA cut-through proxy.
- In the **auth access-list** command, permit user NONE rules guarantee only unauthenticated trigger cut-through proxy. Ideally, they should be the last lines.

```
hostname(config)# access-list listenerAuth extended permit tcp any any
hostname(config)# aaa authentication match listenerAuth inside ldap
hostname(config)# aaa authentication listener http inside port 8888
```

```
hostname(config)# access-list 100 ex permit ip user SAMPLE\user1 any any
hostname(config)# access-list 100 ex deny ip user SAMPLE\user2 any any
hostname(config)# access-list 100 ex permit ip user NONE any any
hostname(config)# access-list 100 ex deny any any
hostname(config)# access-group 100 in interface inside
hostname(config)# aaa authenticate match 200 inside user-identity
```

VPN Filter Example

Some traffic might need to bypass the Identity Firewall.

The ASA reports users logging in through VPN authentication or a web portal (cut-through proxy) to the AD Agent, which distributes the user information to all registered ASA devices. Specifically, the IP-user mapping of authenticated users is forwarded to all ASA contexts that include the input interface where HTTP/HTTPS packets are received and authenticated. The ASA designates users logging in through a VPN as belonging the LOCAL domain.

There are two different ways to apply identity firewall (IDFW) rules to VPN users:

- Apply VPN-Filter with bypassing access-list check disabled
- Apply VPN-Filter with bypassing access-list check enabled

VPN with IDFW Rule -1 Example

By default, the **sysopt connection permit-vpn** command is enabled and VPN traffic is exempted from an access list check. To apply interface-based ACL rules for VPN traffic, VPN traffic access list bypassing needs to be disabled.

In this example, if the user logs in from the outside interface, the IDFW rules control which network resources are accessible. All VPN users are to be stored under the LOCAL domain. Therefore, it is only meaningful to apply the rules for LOCAL users or object groups that include LOCAL users.

```
! Apply VPN-Filter with bypassing access-list check disabled
no sysopt connection permit-vpn
access-list v1 extended deny ip user LOCAL\idfw any 10.0.0.0 255.255.255.0
access-list v1 extended permit ip user LOCAL\idfw any 20.0.0.0 255.255.255.0
access-group v1 in interface outside
```

VPN with IDFW Rule -2 Example

By default, the **sysopt connection permit-vpn** command is enabled, with VPN traffic access bypassing enabled. A VPN filter can be used to apply the IDFW rules to the VPN traffic. A VPN filter with IDFW rules can be defined in the CLI username and group policy.

In the example, when user idfw logs in, the user can access network resources in the 10.0.00/24 subnet. However, when user user1 logs in, access to network resources in 10.0.00/24 subnet is denied. Note that all VPN users are stored under the LOCAL domain. Therefore, it is only meaningful to apply the rules for LOCAL users or object groups that include LOCAL users.



Note

IDFW rules can only be applied to VPN filters under group policy and are not available in all of the other group policy features.

```
! Apply VPN-Filter with bypassing access-list check enabled
```

```
sysopt connection permit-vpn
access-list v1 extended permit ip user LOCAL\idfw any 10.0.0.0 255.255.255.0
access-list v2 extended deny ip user LOCAL\user1 any 10.0.0.0 255.255.255.0
username user1 password QkBIIVi6IFLEsYv encrypted privilege 0 username user1 attributes
    vpn-group-policy group1 vpn-filter value v2
username idfw password eEm2dmjMaopcGozT encrypted
username idfw attributes
    vpn-group-policy testgroup vpn-filter value v1

sysopt connection permit-vpn
access-list v1 extended permit ip user LOCAL\idfw any 10.0.0.0 255.255.255.0 access-list
v1 extended deny ip user LOCAL\user1 any 10.0.0.0 255.255.255.0 group-policy group1
internal
group-policy group1 attributes

    vpn-filter value v1
vpn-tunnel-protocol ikev1 l2tp-ipsec ssl-client ssl-clientless
```

Monitoring the Identity Firewall

This section includes the following topics:

- [Monitoring AD Agents, page 38-23](#)
- [Monitoring Groups, page 38-23](#)
- [Monitoring Memory Usage for the Identity Firewall, page 38-23](#)
- [Monitoring Users for the Identity Firewall, page 38-24](#)

Monitoring AD Agents

To obtain troubleshooting information for the AD Agent, use one of the following commands:

- **show user-identity ad-agent**
- **show user-identity ad-agent statistics**

These commands display 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 obtain troubleshooting information for the user groups configured for the Identity Firewall, use the **show user-identity group** command.

Monitoring Memory Usage for the Identity Firewall

To obtain troubleshooting information for memory usage for the Identity Firewall, use the **show user-identity memory** command.

The command 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 38-14.

Monitoring Users for the Identity Firewall

To obtain troubleshooting information for the AD Agent, enter one of the following commands:

- **show user-identity user all list**
- **show user-identity user active user *domain\user-name* list detail**

These commands display the following information for users:

```
domain\user_name   Status (active or inactive)   Connections           Minutes Idle
```

```
domain\user_name   Active 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 command **user-identity action domain-controller-down *domain_name* disable-user-identity-rule** is configured and the specified domain is down, or if the **user-identity action ad-agent-down disable-user-identity-rule** command is configured and the AD Agent is down, all the logged-in users have the disabled status.

Feature History for the Identity Firewall

Table 38-1 lists the release history for this feature.

Table 38-1 Feature History for the Identity Firewall

Feature Name	Releases	Feature Information
Identity Firewall	8.4(2)	<p>The Identity Firewall feature was introduced.</p> <p>We introduced or modified the following commands: user-identity enable, user-identity default-domain, user-identity domain, user-identity logout-probe, user-identity inactive-user-timer, user-identity poll-import-user-group-timer, user-identity action netbios-response-fail, user-identity user-not-found, user-identity action ad-agent-down, user-identity action mac-address-mismatch, user-identity action domain-controller-down, user-identity ad-agent active-user-database, user-identity ad-agent hello-timer, user-identity ad-agent aaa-server, user-identity update import-user, user-identity static user, dns domain-lookup, dns poll-timer, dns expire-entry-timer, object-group user, show user-identity, show dns, clear configure user-identity, clear dns, debug user-identity.</p>



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](#)
- [Configuration Example, page 39-25](#)
- [Monitoring Cisco TrustSec, page 39-25](#)
- [Feature History for the Cisco TrustSec Integration, page 39-26](#)

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 a 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 as 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

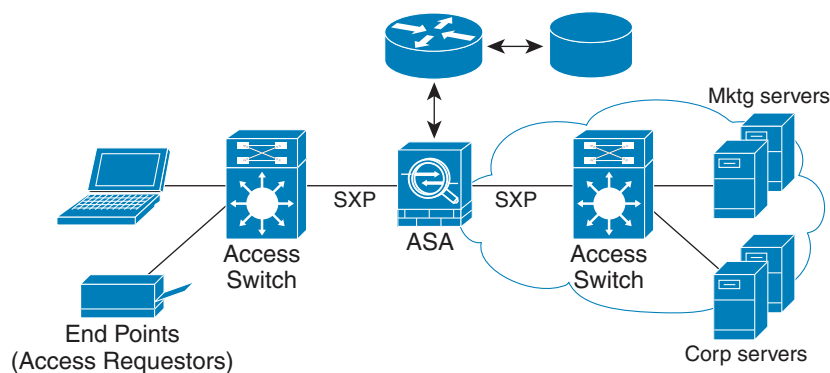
Security policy enforcement is based on security group name. An end-point device attempts to access a resource in the data center. Compared to traditional IP-based policies configured on firewalls, identity-based policies are configured based on user and device identities. For example, mktg-contractor 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



304015

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 39-17.

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 security 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 TrustSec 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 39-17](#).

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 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 39-17.
- 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 39-23](#) 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 39-17](#)
- [Configuring the Security Exchange Protocol \(SXP\), page 39-19](#)
- [Adding an SXP Connection Peer, page 39-22](#)
- [Refreshing Environment Data, page 39-23](#)
- [Configuring the Security Policy, page 39-23](#)

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 39-17.
- Step 3** Enable and set the default values for SXP.
See the [“Configuring the Security Exchange Protocol \(SXP\)”](#) section on page 39-19.
- Step 4** Add SXP connection peers for the Cisco TrustSec architecture.
See the [“Adding an SXP Connection Peer”](#) section on page 39-22.
- Step 5** As necessary, refresh environment data for the ASA.
See the [“Refreshing Environment Data”](#) section on page 39-23.
- Step 6** Configure the security policy.
See the [“Configuring the Security Policy”](#) section on page 39-23.
-

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 AAA server group for the ISE on the ASA, perform the following steps:

	Command	Purpose
Step 1	<pre>hostname(config)# aaa-server <i>server-tag</i> protocol radius</pre> <p>Example: <pre>hostname(config)# aaa-server ISEserver protocol radius</pre></p>	<p>Creates the AAA server group and configures the AAA server parameters for the ASA to communicate with the ISE server.</p> <p><i>server-tag</i> specifies the server group name.</p> <p>See the “Creating a Security Group on the ISE” section on page 39-12 for more information.</p>
Step 2	<pre>hostname(config-aaa-server-group)# exit</pre>	Exits from the AAA server group configuration mode.
Step 3	<pre>hostname(config)# aaa-server <i>server-tag</i> (<i>interface-name</i>) host <i>server-ip</i></pre> <p>Example: <pre>hostname(config)# aaa-server ISEserver (inside) host 192.0.2.1</pre></p>	<p>Configures a AAA server as part of a AAA server group and sets host-specific connection data.</p> <p><i>interface-name</i> specifies the network interface where the ISE server resides. The parentheses are required in this parameter.</p> <p><i>server-tag</i> is the name of the AAA server group.</p> <p><i>server-ip</i> specifies the IP address of the ISE server.</p>
Step 4	<pre>hostname(config-aaa-server-host)# key <i>key</i></pre> <p>Example: <pre>hostname(config-aaa-server-host)# key myexclusivemumblekey</pre></p>	<p>Specifies the server secret value used to authenticate the ASA with the ISE server.</p> <p><i>key</i> is an alphanumeric keyword up to 127 characters long.</p> <p>If the ISE is also used for user authentication, enter the shared secret that was entered on the ISE when you registered the ASA with the ISE.</p> <p>See the “Registering the ASA with the ISE” section on page 39-11 for more information.</p>
Step 5	<pre>hostname(config-aaa-server-host)# exit</pre>	Exits from the AAA server host configuration mode.
Step 6	<pre>hostname(config)# cts server-group <i>AAA-server-group-name</i></pre> <p>Example: <pre>hostname(config)# cts server-group ISEserver</pre></p>	<p>Identifies the AAA server group that is used by Cisco TrustSec for environment data retrieval.</p> <p><i>AAA-server-group-name</i> is the name of the AAA server group that you specified in Step 1 in the <i>server-tag</i> argument.</p> <p>Only one instance of the server group can be configured on the ASA for Cisco TrustSec.</p>

Examples

The following example shows how to configure the ASA to communicate with the ISE server for Cisco TrustSec integration:

```
hostname(config)# aaa-server ISEserver protocol radius
hostname(config-aaa-server-group)# exit
hostname(config)# aaa-server ISEserver (inside) host 192.0.2.1
hostname(config-aaa-server-host)# key myexclusivemumblekey
hostname(config-aaa-server-host)# exit
hostname(config)# cts server-group ISEserver
```

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.

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, enter the following command:

Command	Purpose
<pre>ciscoasa(config)# cts import-pac filepath password value</pre> <p>Example:</p> <pre>ciscoasa(config)# cts import-pac disk0:/xyz.pac password IDFW-pac99</pre>	<p>Imports a Cisco TrustSec PAC file.</p> <p><i>filepath</i> is entered as one of the following exec mode commands and options:</p> <p>Single Mode</p> <ul style="list-style-type: none"> • disk0: Path and filename on disk0 • disk1: Path and filename on disk1 • flash: Path and filename on flash • ftp: Path and filename on FTP • http: Path and filename on HTTP • https: Path and filename on HTTPS • smb: Path and filename on SMB • tftp: Path and filename on TFTP <p>Multi-mode</p> <ul style="list-style-type: none"> • http: Path and filename on HTTP • https: Path and filename on HTTPS • smb: Path and filename on SMB • tftp: Path and filename on TFTP <p><i>value</i> specifies 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.</p>

Examples

The following example shows how to import a PAC file into the ASA:

```
ciscoasa(config)# cts import pac disk0:/pac123.pac password hideme
PAC file successfully imported
```

The following example shows how to use the terminal to import a PAC file into the ASA:

```
ciscoasa(config)# cts import-pac terminal password A9875Za551
Enter the PAC file data in ASCII hex format
End with the word "quit" on a line by itself.
ciscoasa(exec_pac_hex)# 01002904050000010000000000000000
ciscoasa(exec_pac_hex)# 00000000000000011111111111111111
ciscoasa(exec_pac_hex)# 11111111111111112222222222222222
ciscoasa(exec_pac_hex)# 222222222222222276d7d64b6be4804b
ciscoasa(exec_pac_hex)# 0b4fdca3aeed11950ecd0e47c34157e5
ciscoasa(exec_pac_hex)# 25f4964ed75835cde0adb7e198e0bcdb
ciscoasa(exec_pac_hex)# 6aa8e363b0e4f9b4ac241be9ab576d0b
ciscoasa(exec_pac_hex)# a1fcd34e5dd05dbe1312cbfea072fdb9
ciscoasa(exec_pac_hex)# ee356fb61fe987d2d8f0ac3ef0467627
ciscoasa(exec_pac_hex)# 7f8b137da2b840e16da520468b039bae
ciscoasa(exec_pac_hex)# 36a4d844acc85cdefd7cb2cc58787590
ciscoasa(exec_pac_hex)# ef123882a69b6c37bdbc9320e403024f
```



```
ciscoasa(exec_pac_hex) # 354d42f404ec2d67ef3606575014584b
ciscoasa(exec_pac_hex) # 2796e65ccd6e6c8d14d92448a8b24f6e
ciscoasa(exec_pac_hex) # 47015a21f4f66cf6129d352bdfd4520f
ciscoasa(exec_pac_hex) # 3f0c6f340a80715df4498956efe15dec
ciscoasa(exec_pac_hex) # c08bb9a58cb6cb83ac91a3c40ce61de0
ciscoasa(exec_pac_hex) # 284b743e52fd68e848685e2d78c33633
ciscoasa(exec_pac_hex) # f2b4c5824138fc7bac9d9b83ac58ff9f
ciscoasa(exec_pac_hex) # 1dbc84c416322f1f3c5951cf2132994a
ciscoasa(exec_pac_hex) # a7cf20409df1d0d6621eba2b3af83252
ciscoasa(exec_pac_hex) # 70d0130650122bdb13a83b2dae55533a
ciscoasa(exec_pac_hex) # 4a394f21b441e164
ciscoasa(exec_pac_hex) # quit
PAC Imported Successfully
ciscoasa(config) #
```

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 SXP, perform the following steps:

	Command	Purpose
Step 1	<code>hostname(config)# cts sxp enable</code>	If necessary, enables SXP on the ASA. By default, SXP is disabled. In multi-context mode, you enable SXP in the user context.
Step 2	<code>hostname(config)# cts sxp default source-ip <i>ipaddress</i></code> Example: <code>hostname(config)# cts sxp default source-ip 192.168.1.100</code>	Configures the default source IP address for SXP connections. <i>ipaddress</i> is an IPv4 or IPv6 address. When you configure a default source IP address for SXP connections, you must specify the same address as the ASA outbound interface. If the source IP address does not match the address of the outbound interface, SXP connections fail. When a source IP address for an SXP connection is not configured, the ASA performs a route/ARP lookup to determine the outbound interface for the SXP connection. See the “Adding an SXP Connection Peer” section on page 39-22 for more information.
Step 3	<code>hostname(config)# cts sxp default password [0 8] <i>password</i></code> Example: <code>hostname(config)# cts sxp default password 8 IDFW-TrustSec-99</code>	Configures the default password for TCP MD5 authentication with SXP peers. By default, SXP connections do not have a password set. Configuring an encryption level for the password is optional. If you configure an encryption level, you can only set one level: <ul style="list-style-type: none"> • Level 0—unencrypted cleartext • Level 8—encrypted text <i>password</i> specifies an encrypted string up to 162 characters or an ASCII key string up to 80 characters.

	Command	Purpose
Step 4	<pre>hostname(config)# cts sxp retry period timervalue</pre> <p>Example:</p> <pre>hostname(config)# cts sxp retry period 60</pre>	<p>Specifies the default time interval between ASA attempts to set up new SXP connections between SXP peers. The ASA continues to make connection attempts 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.</p> <p><i>timervalue</i> is in the range of 0 to 64000 seconds. By default, the <i>timervalue</i> is 120 seconds.</p> <p>If you specify 0 seconds, the timer never expires and the ASA does not try to connect to SXP peers.</p> <p>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.</p> <p>We recommend that you configure the retry timer to a different value from its SXP peer devices.</p>
Step 5	<pre>hostname(config)# cts sxp reconciliation period timervalue</pre> <p>Example:</p> <pre>hostname(config)# cts sxp reconciliation period 60</pre>	<p>Specifies the value of the default reconcile timer. After an SXP peer terminates its SXP connection, the ASA starts a hold-down timer.</p> <p>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.</p> <p>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.</p> <p><i>timervalue</i> is in the range of 1 to 64000 seconds. By default, the <i>timervalue</i> is 120 seconds.</p> <p>You cannot specify 0 seconds for the timer, because this value prevents the reconcile timer from starting. Not allowing the reconcile timer to run would keep stale entries for an undefined time and cause unexpected results from policy enforcement.</p>

Examples

The following example shows how to set default values for SXP:

```
hostname(config)# cts sxp enable
hostname(config)# cts sxp default source-ip 192.168.1.100
hostname(config)# cts sxp default password 8 *****
hostname(config)# cts sxp retry period 60
hostname(config)# cts sxp reconcile period 60
```

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:

	Command	Purpose
Step 1	<code>hostname(config)# cts sxp enable</code>	If necessary, enables SXP on the ASA. By default, SXP is disabled.
Step 2	<pre>hostname(config)# cts sxp connection peer peer_ip_address [source source_ip_address] password {default none} [mode {local peer}] {speaker listener}</pre> <p>Example:</p> <pre>ciscoasa(config)# cts sxp connection peer 192.168.1.100 password default mode peer speaker</pre>	<p>Sets up an SXP connection to an SXP peer. SXP connections are set per IP address; a single device pair can service multiple SXP connections.</p> <p><i>peer_ip_address</i> is the IPv4 or IPv6 address of the SXP peer. The peer IP address must be reachable from the ASA outgoing interface.</p> <p><i>source_ip_address</i> is the local IPv4 or IPv6 address of the SXP connection. The source IP address must be the same as the ASA outbound interface or the connection fails.</p> <p>We recommend that you do not configure a source IP address for an SXP connection and allow the ASA to perform a route/ARP lookup to determine the source IP address for the SXP connection.</p> <p>Specifies whether or not to use the authentication key for the SXP connection:</p> <ul style="list-style-type: none"> • default—Use the default password configured for SXP connections. See the “Configuring the Security Exchange Protocol (SXP)” section on page 39-19. • none—Do not use a password for the SXP connection. <p>Specifies the mode of the SXP connection:</p> <ul style="list-style-type: none"> • local—Use the local SXP device. • peer—Use the peer SXP device. <p>Specifies whether the ASA functions as a Speaker or Listener for the SXP connection. See the “About Speaker and Listener Roles on the ASA” section on page 39-6.</p> <ul style="list-style-type: none"> • speaker—The ASA can forward IP-SGT mapping to upstream devices. • listener—The ASA can receive IP-SGT mapping from downstream devices.

Examples

The following example shows how to configure SXP peers on the ASA:

```
hostname(config)# cts sxp enable
hostname(config)# cts sxp connection peer 192.168.1.100 password default mode peer speaker
hostname(config)# cts sxp connection peer 192.168.1.101 password default mode peer
hostname(config)# no cts sxp connection peer 192.168.1.100
hostname(config)# cts sxp connection peer 192.168.1.100 source 192.168.1.1 password default mode peer speaker
hostname(config)# no cts sxp connection peer 192.168.1.100 source 192.168.1.1 password default mode peer
speaker
```

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, enter the following command:

Command	Purpose
<code>cts refresh environment-data</code>	Refreshes the environment data from the ISE and resets the reconcile timer to the configured default value.
Example: <code>ciscoasa(config)# cts refresh environment-data</code>	

Configuring the Security Policy

You can incorporate TrustSec policy in many ASA features. Any feature that uses extended ACLs (unless listed in this chapter as unsupported) can take advantage of TrustSec. You can now add security group arguments to extended ACLs, as well as traditional network-based parameters.

- To configure an extended ACL, see [Chapter 19, “Adding an Extended Access Control List.”](#)
- To configure security group object groups, which can be used in the ACL, see the [“Configuring Local User Groups”](#) section on page 17-11.

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.

Examples

The following example shows how to create an ACL that uses a locally defined security object group:

```
object-group security objgrp-it-admin
  security-group name it-admin-sg-name
  security-group tag 1
object-group security objgrp-hr-admin
  security-group name hr-admin-sg-name // single sg_name
  group-object it-admin // locally defined object-group as nested object
object-group security objgrp-hr-servers
  security-group name hr-servers-sg-name
object-group security objgrp-hr-network
  security-group tag 2
access-list hr-acl permit ip object-group-security objgrp-hr-admin any
object-group-security objgrp-hr-servers
```

The ACL configured in the previous example can be activated by configuring an access group or the Modular Policy Framework.

Additional examples:

```
!match src hr-admin-sg-name from any network to dst host 172.23.59.53
  access-list idw-acl permit ip security-group name hr-admin-sg-name any host 172.23.59.53
!match src hr-admin-sg-name from host 10.1.1.1 to dst any
  access-list idfw-acl permit ip security-group name hr-admin-sg-name host 10.1.1.1 any
!match src tag 22 from any network to dst hr-servers-sg-name any network
  access-list idfw-acl permit ip security-group tag 22 any security-group name hr-servers-sg-name any
!match src user mary from any host to dst hr-servers-sg-name any network
  access-list idfw-acl permit ip user CSC0\mary any security-group name hr-servers-sg-name any
!match src objgrp-hr-admin from any network to dst objgrp-hr-servers any network
  access-list idfw-acl permit ip object-group-security objgrp-hr-admin any object-group-security
  objgrp-hr-servers any
!match src user Jack from objgrp-hr-network and ip subnet 10.1.1.0/24 to dst objgrp-hr-servers any network
  access-list idfw-acl permit ip user CSC0\Jack object-group-security objgrp-hr-network 10.1.1.0
  255.255.255.0 object-group-security objgrp-hr-servers any
!match src user Tom from security-group mktg any google.com
object network net-google
  fqdn google.com
  access-list sgacl permit ip sec name mktg any object net-google
! If user Tom or object_group security objgrp-hr-admin needs to be matched, multiple ACEs can be defined as
follows:
  access-list idfw-acl2 permit ip user CSC0\Tom 10.1.1.0 255.255.255.0 object-group-security
  objgrp-hr-servers any
  access-list idfw-acl2 permit ip object-group-security objgrp-hr-admin 10.1.1.0 255.255.255.0
  object-group-security objgrp-hr-servers any
```

Configuration Example

The following example shows how to configure the ASA to use Cisco TrustSec:

```
// Import an encrypted CTS PAC file
cts import-pac asa.pac password Cisco
// Configure ISE for environment data download
aaa-server cts-server-list protocol radius
aaa-server cts-server-list host 10.1.1.100 cisco123
cts server-group cts-server-list
// Configure SXP peers
cts sxp enable
cts sxp connection peer 192.168.1.100 password default mode peer speaker
//Configure security-group based policies
object-group security objgrp-it-admin
  security-group name it-admin-sg-name
  security-group tag 1
object-group security objgrp-hr-admin
  security-group name hr-admin-sg-name
  group-object it-admin
object-group security objgrp-hr-servers
  security-group name hr-servers-sg-name
access-list hr-acl permit ip object-group-security objgrp-hr-admin any
object-group-security objgrp-hr-servers
```

Monitoring Cisco TrustSec

To monitor Cisco TrustSec on the ASA, enter one or more of the following commands:

Command	Purpose
show running-config cts	Displays the configured default values for the Cisco TrustSec infrastructure and the SXP commands.
show cts sxp connections	Displays the SXP connections on the ASA for a particular user context when multi-context mode is used.
show conn security-group	Displays data for all SXP connections.
show cts environment-data	Displays the Cisco TrustSec environment information contained in the security group table on the ASA.
show cts sgt-map	Displays the IP address-security group table manager entries in the control path.
show asp table cts sgt-map	Displays the IP address-security group table mapping entries from the IP address-security group table mapping database maintained in the datapath.
show cts 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.

Table 39-3 Feature History for the Cisco TrustSec Integration

Feature Name	Platform Releases	Feature Information
Cisco TrustSec Integration	9.0(1)	<p>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.</p> <p>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.</p> <p>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.</p> <p>We introduced or modified the following commands: access-list extended, cts sxp enable, cts server-group, cts sxp default, cts sxp retry period, cts sxp reconciliation period, cts sxp connection peer, cts import-pac, cts refresh environment-data, object-group security, security-group, show running-config cts, show running-config object-group, clear configure cts, clear configure object-group, show cts pac, show cts environment-data, show cts environment-data sg-table, show cts sxp connections, show object-group, show configure security-group, clear cts environment-data, debug cts, packet-tracer.</p>



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 40-8](#)
- [Prerequisites for Local Certificates, page 40-9](#)
- [Guidelines and Limitations, page 40-9](#)
- [Configuring Digital Certificates, page 40-10](#)
- [Monitoring Digital Certificates, page 40-43](#)
- [Feature History for Certificate Management, page 40-45](#)

Information About Digital Certificates

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 40-2](#)
- [Certificate Scalability, page 40-2](#)
- [Key Pairs, page 40-2](#)
- [Trustpoints, page 40-3](#)
- [Revocation Checking, page 40-4](#)
- [The Local CA, page 40-6](#)
- [Using Certificates and User Login Credentials, page 40-7](#)

Public Key Cryptography

Digital signatures, enabled by public key cryptography, provide a way to authenticate devices and users. In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other.

In simple terms, a signature is formed when data is encrypted with a private key. The signature is attached to the data and sent to the receiver. The receiver applies the public key of the sender to the data. If the signature sent with the data matches the result of applying the public key to the data, the validity of the message is established.

This process relies on the receiver having a copy of the public key of the sender and a high degree of certainty that this key belongs to the sender, not to someone pretending to be the sender.

Obtaining the public key of a sender is normally handled externally or through an operation performed at installation. For example, most web browsers are configured with the root certificates of several CAs by default. For VPN, the IKE protocol, a component of IPsec, can use digital signatures to authenticate peer devices before setting up security associations.

Certificate Scalability

Without digital certificates, you must manually configure each IPsec peer for each peer with which it communicates; as a result, each new peer that you add to a network would require a configuration change on each peer with which it needs to communicate securely.

When you use digital certificates, each peer is enrolled with a CA. When two peers try to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new peer is added to the network, you enroll that peer with a CA and none of the other peers need modification. When the new peer attempts an IPsec connection, certificates are automatically exchanged and the peer can be authenticated.

With a CA, a peer authenticates itself to the remote peer by sending a certificate to the remote peer and performing some public key cryptography. Each peer sends its unique certificate, which was issued by the CA. This process works because each certificate encapsulates the public key for the associated peer, each certificate is authenticated by the CA, and all participating peers recognize the CA as an authenticating authority. The process is called IKE with an RSA signature.

The peer can continue sending its certificate for multiple IPsec sessions, and to multiple IPsec peers, until the certificate expires. When its certificate expires, the peer administrator must obtain a new one from the CA.

CAs can also revoke certificates for peers that no longer participate in IPsec. Revoked certificates are not recognized as valid by other peers. Revoked certificates are listed in a CRL, which each peer may check before accepting a certificate from another peer.

Some CAs have an RA as part of their implementation. An RA is a server that acts as a proxy for the CA, so that CA functions can continue when the CA is unavailable.

Key Pairs

Key pairs are RSA keys, which have the following characteristics:

- RSA keys can be used for SSH or SSL.
- SCEP enrollment supports the certification of RSA keys.

- For the purposes of generating keys, the maximum key modulus for RSA keys is 2048 bits. The default size is 1024. Many SSL connections using identity certificates with RSA key pairs that exceed 1024 bits can cause a high CPU usage on the ASA and rejected clientless logins.
- For signature operations, the supported maximum key size is 4096 bits. 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 **revocation-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.

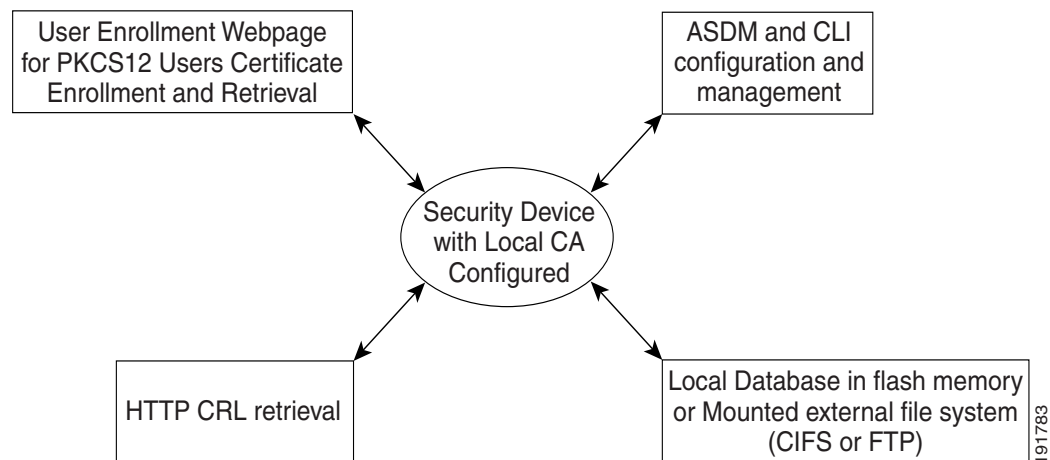
No limits exist on the number of users that can be stored in the local CA user database; however, if flash memory storage issues arise, syslogs are generated to alert the administrator to take action, and the local CA could be disabled until the storage issues are resolved. Flash memory can store a database 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 40-7](#)
- [Using Certificates, page 40-8](#)

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 13-1.
- 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 13-4.

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"phÖV°3é%b0 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 Key Pairs, page 40-11](#)

- [Removing Key Pairs, page 40-12](#)
- [Configuring Trustpoints, page 40-12](#)
- [Configuring CRLs for a Trustpoint, page 40-15](#)
- [Exporting a Trustpoint Configuration, page 40-17](#)
- [Importing a Trustpoint Configuration, page 40-18](#)
- [Configuring CA Certificate Map Rules, page 40-19](#)
- [Obtaining Certificates Manually, page 40-20](#)
- [Obtaining Certificates Automatically with SCEP, page 40-22](#)
- [Configuring Proxy Support for SCEP Requests, page 40-23](#)
- [Enabling the Local CA Server, page 40-24](#)
- [Configuring the Local CA Server, page 40-25](#)
- [Customizing the Local CA Server, page 40-27](#)
- [Debugging the Local CA Server, page 40-28](#)
- [Disabling the Local CA Server, page 40-28](#)
- [Deleting the Local CA Server, page 40-28](#)
- [Configuring Local CA Certificate Characteristics, page 40-29](#)

Configuring Key Pairs

To generate key pairs, perform the following steps:

	Command	Purpose
Step 1	<p>crypto key generate rsa</p> <p>Example: ciscoasa/contexta(config)# crypto key generate rsa</p>	<p>Generates one, general-purpose RSA key pair. The default key modulus is 1024. To specify other modulus sizes, use the modulus keyword.</p> <p>Note Many SSL connections using identity certificates with RSA key pairs that exceed 1024 bits can cause high CPU usage on the ASA and rejected clientless logins.</p>
Step 2	<p>crypto key generate rsa label key-pair-label</p> <p>Example: ciscoasa/contexta(config)# crypto key generate rsa label exchange</p>	<p>(Optional) Assigns a label to each key pair. The label is referenced by the trustpoint that uses the key pair. If you do not assign a label, the key pair is automatically labeled, <i>Default-RSA-Key</i>.</p>

	Command	Purpose
Step 3	show crypto key <i>name of key</i> Example: ciscoasa/contexta(config)# show crypto key examplekey	Verifies key pairs that you have generated.
Step 4	write memory Example: ciscoasa(config)# write memory	Saves the key pair that you have generated.

Removing Key Pairs

To remove key pairs, perform the following steps:

Command	Purpose
crypto key zeroize rsa Example: ciscoasa(config)# crypto key zeroize rsa	Removes key pairs.

Examples

The following example shows how to remove key pairs:

```
ciscoasa(config)# crypto key zeroize rsa
WARNING: All RSA keys will be removed.
WARNING: All device certs issued using these keys will also be removed.

Do you really want to remove these keys? [yes/no] y
```

Configuring Trustpoints

To configure a trustpoint, perform the following steps:

	Command	Purpose
Step 1	crypto ca trustpoint <i>trustpoint-name</i> Example: ciscoasa/contexta(config)# crypto ca trustpoint Main	Creates a trustpoint that corresponds to the CA from which the ASA needs to receive a certificate. Enters the crypto ca trustpoint configuration mode, which controls CA-specific trustpoint parameters that you may configure starting in Step 3. Note When you try to connect, a warning occurs to indicate that the trustpoint does not contain an ID certificate when an attempt is made to retrieve the ID certificate from the trustpoint.
Step 2	Choose one of the following options:	

	Command	Purpose
	enrollment url url Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# enrollment url http://10.29.67.142:80/certsrv/mscep/mscep.dll</pre>	Requests automatic enrollment using SCEP with the specified trustpoint and configures the enrollment URL.
	enrollment terminal Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# enrollment terminal</pre>	Requests manual enrollment with the specified trustpoint by pasting the certificate received from the CA into the terminal.
Step 3	revocation-check crl none <pre>revocation-check crl revocation-check none</pre> Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# revocation-check crl none ciscoasa/contexta(config-ca-trustpoint)# revocation-check crl ciscoasa/contexta(config-ca-trustpoint)# revocation-check none</pre>	Specifies the available CRL configuration options. Note To enable either required or optional CRL checking, make sure that you configure the trustpoint for CRL management after obtaining certificates.
Step 4	crl configure Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# crl configure</pre>	Enters crl configuration mode.
Step 5	email address Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# email example.com</pre>	During enrollment, asks the CA to include the specified e-mail address in the Subject Alternative Name extension of the certificate.
Step 6	enrollment retry period Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# enrollment retry period 5</pre>	(Optional) Specifies a retry period in minutes, and applies only to SCEP enrollment.
Step 7	enrollment retry count Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# enrollment retry period 2</pre>	(Optional) Specifies a maximum number of permitted retries, and applies only to SCEP enrollment.
Step 8	fqdn fqdn Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# fqdn example.com</pre>	During enrollment, asks the CA to include the specified fully qualified domain name in the Subject Alternative Name extension of the certificate.

	Command	Purpose
Step 9	<p>ip-address <i>ip-address</i></p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# ip-address 10.10.100.1</p>	During enrollment, asks the CA to include the IP address of the ASA in the certificate.
Step 10	<p>keypair <i>name</i></p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# keypair exchange</p>	Specifies the key pair whose public key is to be certified.
Step 11	<p>match certificate map-name override ocsp</p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# match certificate examplemap override ocsp</p>	Configures OCSP URL overrides and trustpoints to use for validating OCSP responder certificates.
Step 12	<p>ocsp disable-nonce</p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# ocsp disable-nonce</p>	Disables the nonce extension on an OCSP request. The nonce extension cryptographically binds requests with responses to avoid replay attacks.
Step 13	<p>ocsp url</p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# ocsp url</p>	Configures an OCSP server for the ASA to use to check all certificates associated with a trustpoint rather than the server specified in the AIA extension of the client certificate.
Step 14	<p>password <i>string</i></p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# password mypassword</p>	Specifies a challenge phrase that is registered with the CA during enrollment. The CA usually uses this phrase to authenticate a subsequent revocation request.
Step 15	<p>revocation check</p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# revocation check</p>	Sets one or more methods for revocation checking: CRL, OCSP, and none.
Step 16	<p>subject-name <i>X.500 name</i></p> <p>Example: ciscoasa/contexta(config-ca-trustpoint)# myname X.500 exemplename</p>	During enrollment, asks the CA to include the specified subject DN in the certificate. If a DN string includes a comma, enclose the value string within double quotes (for example, O="Company, Inc.>").

	Command	Purpose
Step 17	serial-number Example: <pre>ciscoasa/contexta(config-ca-trustpoint)# serial number JMX1213L2A7</pre>	During enrollment, asks the CA to include the ASA serial number in the certificate.
Step 18	write memory Example: <pre>ciscoasa/contexta(config)# write memory</pre>	Saves the running configuration.

Configuring CRLs for a Trustpoint

To use mandatory or optional CRL checking during certificate authentication, you must configure CRLs for each trustpoint. To configure CRLs for a trustpoint, perform the following steps:

	Command	Purpose
Step 1	crypto ca trustpoint trustpoint-name Example: <pre>ciscoasa (config)# crypto ca trustpoint Main</pre>	Enters crypto ca trustpoint configuration mode for the trustpoint whose CRL configuration you want to modify. Note Make sure that you have enabled CRLs before entering this command. In addition, the CRL must be available for authentication to succeed.
Step 2	crl configure Example: <pre>ciscoasa (config-ca-trustpoint)# crl configure</pre>	Enters crl configuration mode for the current trustpoint. Tip To set all CRL configuration parameters to default values, use the default command. At any time during CRL configuration, reenter this command to restart the procedure.
Step 3	Do one of the following:	
	policy cdp Example: <pre>ciscoasa (config-ca-crl)# policy cdp</pre>	Configures retrieval policy. CRLs are retrieved only from the CRL distribution points specified in authenticated certificates. Note SCEP retrieval is not supported by distribution points specified in certificates. To continue, go to Step 5.
	policy static Example: <pre>ciscoasa (config-ca-crl)# policy static</pre>	Configures retrieval policy. CRLs are retrieved only from URLs that you configure. To continue, go to Step 4.

	Command	Purpose
	<p>policy both</p> <p>Example: ciscoasa (config-ca-crl)# policy both</p>	<p>Configures retrieval policy. CRLs are retrieved from CRL distribution points specified in authenticated certificates and from URLs that you configure.</p> <p>To continue, go to Step 4.</p>
Step 4	<p>url n url</p> <p>Example: ciscoasa (config-ca-crl)# url 2 http://www.example.com</p>	<p>If you used the keywords static or both when you configured the CRL policy, you must configure URLs for CRL retrieval. You can enter up to five URLs, ranked 1 through 5. The <i>n</i> is the rank assigned to the URL. To remove a URL, use the no url n command.</p>
Step 5	<p>protocol http ldap scep</p> <p>Example: ciscoasa (config-ca-crl)# protocol http</p>	<p>Configures the retrieval method. Specifies HTTP, LDAP, or SCEP as the CRL retrieval method.</p>
Step 6	<p>cache-time refresh-time</p> <p>Example: ciscoasa (config-ca-crl)# cache-time 420</p>	<p>Configures how long the ASA caches CRLs for the current trustpoint. <i>refresh-time</i> is the number of minutes that the ASA waits before considering a CRL stale.</p>
Step 7	<p>Do one of the following:</p>	
	<p>enforcenextupdate</p> <p>Example: ciscoasa (config-ca-crl)# enforcenextupdate</p>	<p>Requires the NextUpdate field in CRLs. This is the default setting.</p>
	<p>no enforcenextupdate</p> <p>Example: ciscoasa (config-ca-crl)# no enforcenextupdate</p>	<p>Allows the NextUpdate field to be absent in CRLs.</p>
Step 8	<p>ldap-defaults server</p> <p>Example: ciscoasa (config-ca-crl)# ldap-defaults ldap1</p>	<p>Identifies the LDAP server to the ASA if LDAP is specified as the retrieval protocol. You can specify the server by DNS hostname or by IP address. You can also provide a port number if the server listens for LDAP queries on a port other than the default of 389.</p> <p>Note If you use a hostname instead of an IP address to specify the LDAP server, make sure that you have configured the ASA to use DNS.</p>
Step 9	<p>ldap-dn admin-DN password</p> <p>Example: ciscoasa (config-ca-crl)# ldap-dn cn=admin,ou=devtest,o=engineering c001RunZ</p>	<p>Allows CRL retrieval if the LDAP server requires credentials.</p>

	Command	Purpose
Step 10	<code>crypto ca crl request trustpoint</code> Example: <code>ciscoasa (config-ca-crl)# crypto ca crl request Main</code>	Retrieves the current CRL from the CA represented by the specified trustpoint and tests the CRL configuration for the current trustpoint.
Step 11	<code>write memory</code> Example: <code>ciscoasa (config)# write memory</code>	Saves the running configuration.

Exporting a Trustpoint Configuration

To export a trustpoint configuration, enter the following command:

Command	Purpose
<code>crypto ca export trustpoint</code> Example: <code>ciscoasa(config)# crypto ca export Main</code>	Exports a trustpoint configuration with all associated keys and certificates in PKCS12 format. The ASA displays the PKCS12 data in the terminal. You can copy the data. The trustpoint data is password protected; however, if you save the trustpoint data in a file, make sure that the file is in a secure location.

Examples

The following example exports PKCS12 data for the trustpoint Main with the passphrase Wh0zits:

```
ciscoasa (config)# crypto ca export Main pkcs12 Wh0zits
```

```
Exported pkcs12 follows:
```

```
[ PKCS12 data omitted ]
```

```
---End - This line not part of the pkcs12---
```

Importing a Trustpoint Configuration

To import a trustpoint configuration, enter the following command:

Command	Purpose
crypto ca import trustpoint pkcs12 Example: <pre>ciscoasa(config)# crypto ca import Main pkcs12</pre>	Imports keypairs and issued certificates that are associated with a trustpoint configuration. The ASA prompts you to paste the text into the terminal in base 64 format. The key pair imported with the trustpoint is assigned a label that matches the name of the trustpoint that you create. Note If an ASA has trustpoints that share the same CA, you can use only one of the trustpoints that share the CA to validate user certificates. To control which trustpoint that shares a CA is used for validation of user certificates issued by that CA, use the support-user-cert-validation keyword.

Examples

The following example manually imports PKCS12 data to the trustpoint Main with the passphrase Wh0zits:

```
ciscoasa (config)# crypto ca import Main pkcs12 Wh0zits
```

Enter the base 64 encoded pkcs12.

End with a blank line or the word "quit" on a line by itself:

```
[ PKCS12 data omitted ]
```

```
quit
```

```
INFO: Import PKCS12 operation completed successfully
```

The following example manually imports a certificate for the trustpoint Main:

```
ciscoasa (config)# crypto ca import Main certificate
```

```
% The fully-qualified domain name in the certificate will be:
```

```
securityappliance.example.com
```

Enter the base 64 encoded certificate.

End with a blank line or the word "quit" on a line by itself

```
[ certificate data omitted ]
```

```
quit
```

```
INFO: Certificate successfully imported
```

Configuring CA Certificate Map Rules

You can configure rules based on the Issuer and Subject fields of a certificate. Using the rules you create, you can map IPsec peer certificates to tunnel groups with the **tunnel-group-map** command. The ASA supports one CA certificate map, which can include many rules.

To configure a CA certificate map rule, perform the following steps:

	Command	Purpose
Step 1	crypto ca certificate map <i>sequence-number</i> Example: ciscoasa(config)# crypto ca certificate map 1	Enters CA certificate map configuration mode for the rule you want to configure and specifies the rule index number.
Step 2	issuer-name <i>DN-string</i> Example: ciscoasa(config-ca-cert-map)# issuer-name cn=asa.example.com	Specifies the distinguished name of all issued certificates, which is also the subject-name DN of the self-signed CA certificate. Use commas to separate attribute-value pairs. Insert quotation marks around any value that includes a comma. An issuer-name must be less than 500 alphanumeric characters. The default issuer-name is <i>cn=hostame.domain-name</i> .
Step 3	subject-name attr <i>tag eq co ne nc string</i> Example: ciscoasa(config-ca-cert-map)# subject-name attr cn eq mycert	Specifies tests that the ASA can apply to values found in the Subject field of certificates. The tests can apply to specific attributes or to the entire field. You can configure many tests per rule, and all the tests you specify with these commands must be true for a rule to match a certificate. The following are valid operators: <ul style="list-style-type: none"> • eq—The field or attribute must be identical to the value given. • ne—The field or attribute cannot be identical to the value given. • co—Part or all of the field or attribute must match the value given. • nc—No part of the field or attribute can match the value given.
Step 4	write memory Example: ciscoasa (config)# write memory	Saves the running configuration.

Obtaining Certificates Manually

To obtain certificates manually, perform the following steps:

	Command	Purpose
Step 1	<pre>crypto ca authenticate trustpoint</pre> <p>Example:</p> <pre>ciscoasa(config)# crypto ca authenticate Main Enter the base 64 encoded CA certificate. End with a blank line or the word "quit" on a line by itself MIIDRTCCAu+gAwIBAgIQKVCqP/KW74VP0NZZL+JbRTANBgkqhkiG 9w0BAQUFADCB [certificate data omitted] /7QEM8izy0EOTSErKu7Ng76jwf5e4qtkQ== quit</pre> <pre>INFO: Certificate has the following attributes: Fingerprint: 24b81433 409b3fd5 e5431699 8d490d34 Do you accept this certificate? [yes/no]: y Trustpoint CA certificate accepted.</pre> <pre>% Certificate successfully imported</pre>	<p>Imports the CA certificate for the configured trustpoint.</p> <p>Note This step assumes that you have already obtained a base-64 encoded CA certificate from the CA represented by the trustpoint.</p> <p>Whether a trustpoint requires that you manually obtain certificates is determined by the use of the enrollment terminal command when you configure the trustpoint. For more information, see the “Configuring Trustpoints” section on page 40-12.</p>
Step 2	<pre>crypto ca enroll trustpoint</pre> <p>Example:</p> <pre>ciscoasa(config)# crypto ca enroll Main % Start certificate enrollment ..</pre> <pre>% The fully-qualified domain name in the certificate will be: securityappliance.example.com</pre> <pre>% Include the device serial number in the subject name? [yes/no]: n</pre> <pre>Display Certificate Request to terminal? [yes/no]: y Certificate Request follows:</pre> <pre>MIIBoDCCAQkCAQAwIzEhMB8GCSqGSIb3DQEJAhYSRmVyYWxQaXgu Y2lzY28uY29t [certificate request data omitted] jF4waw68eOxQxVmdgMWeQ+RbIOYmvt8g6hnBTrd0GdqjVLT</pre> <pre>---End - This line not part of the certificate request---</pre> <pre>Redisplay enrollment request? [yes/no]: n</pre>	<p>Enrolls the ASA with the trustpoint. Generates a certificate for signing data and depending on the type of keys that you have configured, for encrypting data.</p> <p>If you use separate RSA keys for signing and encryption, the crypto ca enroll command displays two certificate requests, one for each key. If you use general-purpose RSA keys for both signing and encryption, the crypto ca enroll command displays one certificate request.</p> <p>To complete enrollment, obtain a certificate for all certificate requests generated by the crypto ca enroll command from the CA represented by the applicable trustpoint. Make sure that the certificate is in base-64 format.</p>

	Command	Purpose
Step 3	<p>crypto ca import trustpoint certificate</p> <p>Example: ciscoasa (config)# crypto ca import Main certificate % The fully-qualified domain name in the certificate will be: securityappliance.example.com</p> <p>Enter the base 64 encoded certificate. End with a blank line or the word "quit" on a line by itself [certificate data omitted] quit INFO: Certificate successfully imported</p>	Imports each certificate you receive from the CA. Requests that you paste the certificate to the terminal in base-64 format.
Step 4	<p>show crypto ca server certificate</p> <p>Example: ciscoasa(config)# show crypto ca server certificate Main</p>	Verifies that the enrollment process was successful by displaying certificate details issued for the ASA and the CA certificate for the trustpoint.
Step 5	<p>write memory</p> <p>Example: ciscoasa(config)# write memory</p>	<p>Saves the running configuration.</p> <p>Repeat these steps for each trustpoint that you configure for manual enrollment.</p>

Obtaining Certificates Automatically with SCEP

To obtain certificates automatically using SCEP, perform the following steps:

	Command	Purpose
Step 1	<p><code>crypto ca authenticate trustpoint</code></p> <p>Example: <pre>ciscoasa/contexta(config)# crypto ca authenticate Main</pre></p>	<p>Obtains the CA certificate for the configured trustpoint.</p> <p>Note This step assumes that you have already obtained a base-64 encoded CA certificate from the CA represented by the trustpoint.</p> <p>When you configure the trustpoint, use of the enrollment url command determines whether or not you must obtain certificates automatically via SCEP. For more information, see the “Configuring Trustpoints” section on page 40-12.</p>
Step 2	<p><code>crypto ca enroll trustpoint</code></p> <p>Example: <pre>ciscoasa/contexta(config)# crypto ca enroll Main</pre></p>	<p>Enrolls the ASA with the trustpoint. Retrieves a certificate for signing data and depending on the type of keys that you have configured, for encrypting data. Before entering this command, contact the CA administrator, who may need to authenticate the enrollment request manually before the CA grants certificates.</p> <p>If the ASA does not receive a certificate from the CA within one minute (the default) of sending a certificate request, it resends the certificate request. The ASA continues sending a certificate request each minute until a certificate is received.</p> <p>If the fully qualified domain name configured for the trustpoint is not identical to the fully qualified domain name of the ASA, including the case of the characters, a warning appears. To resolve this issue, exit the enrollment process, make any necessary corrections, and reenter the crypto ca enroll command.</p> <p>Note If the ASA reboots after you have issued the crypto ca enroll command but before you have received the certificate, reenter the crypto ca enroll command and notify the CA administrator.</p>

	Command	Purpose
Step 3	show crypto ca server certificate Example: <pre>ciscoasa/contexta(config)# show crypto ca server certificate Main</pre>	Verifies that the enrollment process was successful by displaying certificate details issued for the ASA and the CA certificate for the trustpoint.
Step 4	write memory Example: <pre>ciscoasa/contexta(config)# write memory</pre>	Saves the running configuration.

Configuring Proxy Support for SCEP Requests

To configure the ASA to authenticate remote access endpoints using third-party CAs, perform the following steps:

	Command	Purpose
Step 1	crypto ikev2 enable outside client-services port portnumber Example: <pre>ciscoasa(config-tunnel-ipsec)# crypto ikev2 enable outside client-services</pre>	Enables client services. Note Needed only if you support IKEv2. Enter this command in tunnel-group ipsec-attributes configuration mode. The default port number is 443.
Step 2	scep-enrollment enable Example: <pre>ciscoasa(config-tunnel-general)# scep-enrollment enable</pre> INFO: 'authentication aaa certificate' must be configured to complete setup of this option.	Enables SCEP enrollment for the tunnel group. Enter this command in tunnel-group general-attributes configuration mode.
Step 3	scep-forwarding-url value URL Example: <pre>ciscoasa(config-group-policy)# scep-forwarding-url value http://ca.example.com:80/</pre>	Enrolls the SCEP CA for the group policy. Enter this command once per group policy to support a third-party digital certificate. Enter the command in group-policy general-attributes configuration mode. <i>URL</i> is the SCEP URL on the CA.
Step 4	secondary-pre-fill-username clientless hide use-common-password password Example: <pre>ciscoasa(config)# tunnel-group remotegrp webvpn-attributes ciscoasa(config-tunnel-webvpn)# secondary-pre-fill-username clientless hide use-common-password secret</pre>	Supplies a common, secondary password when a certificate is unavailable for WebLaunch support of the SCEP proxy. You must use the hide keyword to support the SCEP proxy. For example, a certificate is not available to an endpoint requesting one. Once the endpoint has the certificate, AnyConnect disconnects, then reconnects to the ASA to qualify for a DAP policy that provides access to internal network resources.

	Command	Purpose
Step 5	<pre>secondary-pre-fill-username ssl-client hide use-common-password password</pre> <p>Example:</p> <pre>ciscoasa(config-tunnel-webvpn)# secondary-pre-fill-username ssl-client hide use-common-password secret</pre>	<p>Hides the secondary prefill username for AnyConnect VPN sessions.</p> <p>Despite the ssl-client keyword inherited from earlier releases, use this command to support AnyConnect sessions that use either IKEv2 or SSL.</p> <p>You must use the hide keyword to support the SCEP proxy.</p>
Step 6	<pre>secondary-username-from-certificate {use-entire-name use-script {primary_attr [secondary_attr]}} [no-certificate-fallback cisco-secure-desktop machine-unique-id]</pre> <p>Example:</p> <pre>ciscoasa(config-tunnel-webvpn)# secondary-username-from-certificate CN no-certificate-fallback cisco-secure-desktop machine-unique-id</pre>	<p>Supplies the username when a certificate is unavailable.</p>

Enabling the Local CA Server

Before enabling the local CA server, you must first create a passphrase of at least seven characters to encode and archive a PKCS12 file that includes the local CA certificate and keypair to be generated. The passphrase unlocks the PKCS12 archive if the CA certificate or keypair is lost.

To enable the local CA server, perform the following commands:

	Command	Purpose
Step 1	<pre>crypto ca server</pre> <p>Example:</p> <pre>ciscoasa (config)# crypto ca server</pre>	<p>Enters local ca server configuration mode. Allows you to configure and manage a local CA.</p>
Step 2	<pre>no shutdown</pre> <p>Example:</p> <pre>ciscoasa (config-ca-server)# no shutdown</pre>	<p>Enables the local CA server. Generates the local CA server certificate, keypair and necessary database files, and archives the local CA server certificate and keypair to storage in a PKCS12 file. Requires an 8-65 alphanumeric character password. After initial startup, you can disable the local CA without being prompted for the passphrase.</p> <p>Note After you enable the local CA server, save the configuration to make sure that the local CA certificate and keypair are not lost after a reboot occurs.</p>

Examples

The following example enables the local CA server:

```
hostname (config)# crypto ca server
```



```

ciscoasa (config-ca-server)# no shutdown

% Some server settings cannot be changed after CA certificate generation.
% Please enter a passphrase to protect the private key
% or type Return to exit

Password: caserver

Re-enter password: caserver

Keypair generation process begin. Please wait...

```

The following is sample output that shows local CA server configuration and status:

```

Certificate Server LOCAL-CA-SERVER:
  Status: enabled
  State: enabled
  Server's configuration is locked (enter "shutdown" to unlock it)
  Issuer name: CN=wz5520-1-16
  CA certificate fingerprint/thumbprint: (MD5)
    76dd1439 ac94fdb3 74a0a89f cb815acc
  CA certificate fingerprint/thumbprint: (SHA1)
    58754ffd 9f19f9fd b13b4b02 15b3e4be b70b5a83
  Last certificate issued serial number: 0x6
  CA certificate expiration timer: 14:25:11 UTC Jan 16 2008
  CRL NextUpdate timer: 16:09:55 UTC Jan 24 2007
  Current primary storage dir: flash:

```

Configuring the Local CA Server

To configure the local CA server, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Generates the local CA.
Step 2	smtp from-address <i>e-mail_address</i> Example: ciscoasa (config-ca-server) # smtp from-address SecurityAdmin@example.com	Specifies the SMTP from-address, a valid e-mail address that the local CA uses as a from address when sending e-mail messages that deliver OTPs for an enrollment invitation to users.

	Command	Purpose
Step 3	<p>subject-name-default dn</p> <p>Example: hostname (config-ca-server)# subject-name-default cn=engineer, o=asc systems, c="US"</p>	<p>(Optional) Specifies the subject-name DN that is appended to each username on issued certificates.</p> <p>The subject-name DN and the username combine to form the DN in all user certificates that are issued by the local CA server. If you do not specify a subject-name DN, you must specify the exact subject name DN to be included in a user certificate each time that you add a user to the user database.</p> <p>Note Make sure that you review all optional parameters carefully before you enable the configured local CA, because you cannot change issuer-name and keysize server values after you enable the local CA for the first time.</p>
Step 4	<p>no shutdown</p> <p>Example: hostname (config-ca-server)# no shutdown</p>	<p>Creates the self-signed certificate and associates it with the local CA on the ASA. The self-signed certificate key usage extension has key encryption, key signature, CRL signing, and certificate signing capabilities.</p> <p>Note After the self-signed local CA certificate has been generated, to change any characteristics, you must delete the existing local CA server and completely recreate it.</p> <p>The local CA server keeps track of user certificates, so the administrator can revoke or restore privileges as needed.</p>

Examples

The following example shows how to configure and enable the local CA server using the predefined default values for all required parameters:

```
hostname (config)# crypto ca server
hostname (config-ca-server) # smtp from-address SecurityAdmin@example.com
hostname (config-ca-server) # subject-name-default cn=engineer, o=asc Systems, c=US
hostname (config-ca-server) # no shutdown
```

Customizing the Local CA Server

To configure a customized local CA server, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	issuer-name <i>DN-string</i> Example: hostname (config-ca-server)# issuer-name cn=xx5520,cn=30.132.0.25,ou=DevTest,ou=QA,o=ASC Systems	Specifies parameters that do not have default values.
Step 3	smtp subject <i>subject-line</i> Example: hostname (config-ca-server) # smtp subject Priority E-Mail: Enclosed Confidential Information is Required for Enrollment	Customizes the text that appears in the subject field of all e-mail messages sent from the local CA server
Step 4	smtp from-address <i>e-mail_address</i> Example: hostname (config-ca-server) # smtp from-address SecurityAdmin@example.com	Specifies the e-mail address that is to be used as the From: field of all e-mail messages that are generated by the local CA server.
Step 5	subject-name-default <i>dn</i> Example: hostname (config-ca-server) # subject-name default cn=engineer, o=ASC Systems, c=US	<p>Specifies an optional subject-name DN to be appended to a username on issued certificates. The default subject-name DN becomes part of the username in all user certificates issued by the local CA server.</p> <p>The allowed DN attribute keywords are as follows:</p> <ul style="list-style-type: none"> • C = Country • CN = Common Name • EA = E-mail Address • L = Locality • O = Organization Name • OU = Organization Unit • ST = State/Province • SN = Surname • ST = State/Province <p>Note If you do not specify a subject-name-default to serve as a standard subject-name default, you must specify a DN each time that you add a user.</p>

Debugging the Local CA Server

To debug the newly configured local CA server, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	debug crypto ca server Example: ciscoasa (config-ca-server)# debug crypto ca server	Displays debugging messages when you configure and enable the local CA server. Performs level 1 debugging functions; levels 1-255 are available. Note Debugging commands might slow down traffic on busy networks. Levels 5 and higher are reserved for raw data dumps and should be avoided during normal debugging because of excessive output.

Disabling the Local CA Server

To disable the local CA server, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	shutdown Example: ciscoasa (config-ca-server)# shutdown INFO: Local CA Server has been shutdown.	Disables the local CA server. Disables website enrollment and allows you to modify the local CA server configuration. Stores the current configuration and associated files. After initial startup, you can reenble the local CA without being prompted for the passphrase.

Deleting the Local CA Server

To delete an existing local CA server (either enabled or disabled), enter one of the following commands:

Command	Purpose
Do one of the following:	

Command	Purpose
no crypto ca server Example: ciscoasa (config)# no crypto ca server	Removes an existing local CA server (either enabled or disabled). Note Deleting the local CA server removes the configuration from the ASA. After the configuration has been deleted, it is unrecoverable. Make sure that you also delete the associated local CA server database and configuration files (that is, all files with the wildcard name, LOCAL-CA-SERVER.*).
clear configure crypto ca server Example: ciscoasa (config)# clear config crypto ca server	

Configuring Local CA Certificate Characteristics

You can configure the following characteristics of local CA certificates:

- The name of the certificate issuer as it appears on all user certificates.
- The lifetime of the local CA certificates (server and user) and the CRL.
- The length of the public and private keypairs associated with local CA and user certificates.

This section includes the following topics:

- [Configuring the Issuer Name, page 40-30](#)
- [Configuring the CA Certificate Lifetime, page 40-30](#)
- [Configuring the User Certificate Lifetime, page 40-31](#)
- [Configuring the CRL Lifetime, page 40-32](#)
- [Configuring the Server Keysize, page 40-32](#)
- [Setting Up External Local CA File Storage, page 40-33](#)
- [Downloading CRLs, page 40-35](#)
- [Storing CRLs, page 40-36](#)
- [Setting Up Enrollment Parameters, page 40-37](#)
- [Adding and Enrolling Users, page 40-38](#)
- [Renewing Users, page 40-40](#)
- [Restoring Users, page 40-41](#)
- [Removing Users, page 40-41](#)
- [Revoking Certificates, page 40-42](#)
- [Maintaining the Local CA Certificate Database, page 40-42](#)
- [Rolling Over Local CA Certificates, page 40-42](#)
- [Archiving the Local CA Server Certificate and Keypair, page 40-43](#)

Configuring the Issuer Name

To configure the certificate issuer name, perform the following commands:

	Command	Purpose
Step 1	<code>crypto ca server</code> Example: ciscoasa (config)# <code>crypto ca server</code>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	<code>issuer-name DN-string</code> Example: hostname (config-ca-server)# <code>issuer-name</code> CN=xx5520,CN=30.132.0.25,ou=DevTest,ou=QA,O=ABC Systems	Specifies the local CA certificate subject name. The configured certificate issuer name is both the subject name and issuer name of the self-signed local CA certificate, as well as the issuer name in all issued client certificates and in the issued CRL. The default issuer name in the local CA is in the format, <i>hostname.domainname</i> . Note You cannot change the issuer name value after the local CA is first enabled.

Configuring the CA Certificate Lifetime

To configure the local CA server certificate lifetime, perform the following commands:

	Command	Purpose
Step 1	<code>crypto ca server</code> Example: ciscoasa (config)# <code>crypto ca server</code>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.

	Command	Purpose
Step 2	<p>lifetime ca-certificate time</p> <p>Example: hostname (config-ca-server)# lifetime ca-certificate 365</p>	<p>Determines the expiration date included in the certificate. The default lifetime of a local CA certificate is three years.</p> <p>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.</p>
Step 3	<p>no lifetime ca-certificate</p> <p>Example: hostname (config-ca-server)# no lifetime ca-certificate</p>	<p>(Optional) Resets the local CA certificate lifetime to the default value of three years.</p> <p>The local CA server automatically generates a replacement CA certificate 30 days before it expires, which allows the replacement certificate to be exported and imported onto any other devices for certificate validation of user certificates that have been issued by the local CA certificate after the current local CA certificate has expired. The following preexpiration syslog message is generated:</p> <pre>%ASA-1-717049: Local CA Server certificate is due to expire in days days and a replacement certificate is available for export.</pre> <p>Note When notified of this automatic rollover, the administrator must make sure that the new local CA certificate is imported onto all required devices before it expires.</p>

Configuring the User Certificate Lifetime

To configure the user certificate lifetime, perform the following commands:

	Command	Purpose
Step 1	<p>crypto ca server</p> <p>Example: ciscoasa (config)# crypto ca server</p>	<p>Enters local ca server configuration mode. Allows you to configure and manage a local CA.</p>
Step 2	<p>lifetime certificate time</p> <p>Example: hostname (config-ca-server)# lifetime certificate 60</p>	<p>Sets the length of time that you want user certificates to remain valid.</p> <p>Note Before a user certificate expires, the local CA server automatically initiates certificate renewal processing by granting enrollment privileges to the user several days ahead of the certificate expiration date, setting renewal reminders, and delivering an e-mail message that includes the enrollment username and OTP for certificate renewal. 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.</p>

Configuring the CRL Lifetime

To configure the CRL lifetime, perform the following commands:

	Command	Purpose
Step 1	<pre>crypto ca server</pre> <p>Example: ciscoasa (config)# crypto ca server</p>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	<pre>lifetime crl time</pre> <p>Example: hostname (config-ca-server)# lifetime crl 10</p>	<p>Sets the length of time that you want the CRL to remain valid.</p> <p>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 automatically once each CRL lifetime. If you do not specify a CRL lifetime, the default time period is six hours.</p>
Step 3	<pre>crypto ca server crl issue</pre> <p>Example: ciscoasa(config)# crypto ca server crl issue A new CRL has been issued.</p>	<p>Forces the issuance of a CRL at any time, which immediately updates and regenerates a current CRL to overwrite the existing CRL.</p> <p>Note Do not use this command unless the CRL file has been removed in error or has been corrupted and must be regenerated.</p>

Configuring the Server Keysize

To configure the server keysize, perform the following commands:

	Command	Purpose
Step 1	<pre>crypto ca server</pre> <p>Example: ciscoasa (config)# crypto ca server</p>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	<pre>keysize server</pre> <p>Example: hostname (config-ca-server)# keysize server 2048</p>	<p>Specifies the size of the public and private keys generated at user-certificate enrollment. The keypair size options are 512, 768, 1024, 2048 bits, and the default value is 1024 bits.</p> <p>Note After you have enabled the local CA, you cannot change the local CA keysize, because all issued certificates would be invalidated. To change the local CA keysize, you must delete the current local CA and reconfigure a new one.</p>

Examples

The following is sample output that shows two user certificates in the database.

```

Username: user1
Renewal allowed until: Not Allowed
Number of times user notified: 0
PKCS12 file stored until: 12:45:52 UTC Fri Jan 4 2017
Certificates Issued:
serial:    0x71
issued:   12:45:52 UTC Thu Jan 3 2008
expired:  12:17:37 UTC Sun Dec 31 2017
status:   Not Revoked
Username: user2
Renewal allowed until: Not Allowed
Number of times user notified: 0
PKCS12 file stored until: 12:27:59 UTC Fri Jan 4 2008
Certificates Issued:
serial:    0x2
issued:   12:27:59 UTC Thu Jan 3 2008
expired:  12:17:37 UTC Sun Dec 31 2017
status:   Not Revoked
<--- More --->

```

Setting Up External Local CA File Storage

You can store the local CA server configuration, users, issued certificates, and CRLs in the local CA server database either in flash memory or in an external local CA file system. To configure external local CA file storage, perform the following steps:

	Command	Purpose
Step 1	mount <i>name</i> type Example: hostname (config)# mount mydata type cifs	Accesses configuration mode for the specific file system type.
Step 2	mount <i>name</i> type cifs Example: hostname (config-mount-cifs)# mount mydata type cifs server 10.1.1.10 share myshare domain example.com username user6 password ***** status enable	Mounts a CIFS file system. Note Only the user who mounts a file system can unmount it with the no mount command.

	Command	Purpose
Step 3	<p>crypto ca server</p> <p>Example: ciscoasa (config)# crypto ca server</p>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 4	<p>database path mount-name directory-path</p> <p>Example: hostname (config-ca-server)# database path mydata:newuser</p>	<p>Specifies the location of <i>mydata</i>, the premounted CIFS file system to be used for the local CA server database. Establishes a path to the server and then specifies the local CA file or folder name to use for storage and retrieval. To return local CA file storage to the ASA flash memory, use the no database path command.</p> <p>Note To secure stored local CA files on an external server requires a premounted file system of file type CIFS or FTP that is username-protected and password-protected.</p>
Step 5	<p>write memory</p> <p>Example: ciscoasa (config)# write memory</p>	<p>Saves the running configuration.</p> <p>For external local CA file storage, each time that you save the ASA configuration, user information is saved from the ASA to the premounted file system and file location, <i>mydata:newuser</i>.</p> <p>For flash memory storage, user information is saved automatically to the default location for the start-up configuration.</p>

Examples

The following example shows the list of local CA files that appear in flash memory or in external storage:

```
ciscoasa (config-ca-server)# dir LOCAL* //
Directory of disk0:/LOCAL*

 75  -rwx  32          13:07:49 Jan 20 2007  LOCAL-CA-SERVER.ser
 77  -rwx 229          13:07:49 Jan 20 2007  LOCAL-CA-SERVER.cdb
 69  -rwx   0          01:09:28 Jan 20 2007  LOCAL-CA-SERVER.udb
 81  -rwx 232          19:09:10 Jan 20 2007  LOCAL-CA-SERVER.crl
 72  -rwx 1603         01:09:28 Jan 20 2007  LOCAL-CA-SERVER.p12

127119360 bytes total (79693824 bytes free)
```

Downloading CRLs

To make the CRL available for HTTP download on a given interface or port, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	publish-crl interface interface port portnumber Example: hostname (config-ca-server)# publish-crl outside 70	<p>Opens a port on an interface to make the CRL accessible from that interface. The specified interface and port are used to listen for incoming requests for the CRL. The interface and optional port selections are as follows:</p> <ul style="list-style-type: none"> • inside—Name of interface/GigabitEthernet0/1 • management—Name of interface/Management0/0 • outside—Name of interface/GigabitEthernet0/0 • Port numbers can range from 1-65535. TCP port 80 is the HTTP default port number. <p>Note If you do not specify this command, the CRL is not accessible from the CDP location, because this command is required to open an interface to download the CRL file.</p> <p>The CDP URL can be configured to use the IP address of an interface, and the path of the CDP URL and the filename can also be configured (for example, <code>http://10.10.10.100/user8/my_crl_file</code>).</p> <p>In this case, only the interface with that IP address configured listens for CRL requests, and when a request comes in, the ASA matches the path, <code>/user8/my_crl_file</code> to the configured CDP URL. When the path matches, the ASA returns the stored CRL file.</p> <p>Note The protocol must be HTTP, so the prefix displayed is <code>http://</code>.</p>

Storing CRLs

To establish a specific location for the automatically generated CRL of the local CA, perform the following site-to-site task in either single or multiple context mode:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	cdp-url url Example: ciscoasa(config-ca-server)# cdp-url http://172.16.1.1/pathname/myca.crl	<p>Specifies the CDP to be included in all issued certificates. If you do not configure a specific location for the CDP, the default URL location is <code>http://hostname.domain/+CSCOCA+/asa_ca.crl</code>.</p> <p>The local CA updates and reissues the CRL each time a user certificate is revoked or unrevoked. If no revocation changes occur, the CRL is reissued once each CRL lifetime.</p> <p>If this command is set to serve the CRL directly from the local CA ASA, see the “Downloading CRLs” section on page 40-35 for instructions about opening a port on an interface to make the CRL accessible from that interface.</p> <p>The CRL exists for other devices to validate the revocation of certificates issued by the local CA. In addition, the local CA tracks all issued certificates and status within its own certificate database. Revocation checking is performed when a validating party needs to validate a user certificate by retrieving the revocation status from an external server, which might be the CA that issued the certificate or a server designated by the CA.</p>

Setting Up Enrollment Parameters

To set up enrollment parameters, perform the following commands:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	otp expiration timeout Example: ciscoasa(config-ca-server)# otp expiration 24	Specifies the number of hours that an issued OTP for the local CA enrollment page is valid. The default expiration time is 72 hours. Note The user OTP to enroll for a certificate on the enrollment website is also used as the password to unlock the PKCS12 file that includes the issued certificate and keypair for the specified user.
Step 3	enrollment-retrieval timeout Example: ciscoasa(config-ca-server)# enrollment-retrieval 120	Specifies the number of hours an already-enrolled user can retrieve a PKCS12 enrollment file. This time period begins when the user is successfully enrolled. The default retrieval period is 24 hours. Valid values for the retrieval period range from 1 to 720 hours. The enrollment retrieval period is independent of the OTP expiration period. After the enrollment retrieval time expires, the user certificate and keypair are no longer available. The only way a user may receive a certificate is for the administrator to reinitialize certificate enrollment and allow a user to log in again.

Adding and Enrolling Users

To add a user who is eligible for enrollment in the local CA database, perform the following commands:

	Command	Purpose
Step 1	<pre>crypto ca server user-db add username [dn dn] [email emailaddress]</pre> <p>Example:</p> <pre>hostname (config-ca-server)# crypto ca server user-db add user1 dn user1@example.com, Engineer, Example Company, US, email user1@example.com</pre>	<p>Adds a new user to the local CA database. Options are as follows:</p> <ul style="list-style-type: none"> • <i>username</i>—A string of 4-64 characters, which is the simple username for the user being added. The username can be an e-mail address, which then is used to contact the user as necessary for enrollment invitations. • <i>dn</i>—The distinguished name, a global, authoritative name of an entry in the OSI Directory (X.500) (for example, cn=user1@example.com, cn=Engineer, o=Example Company, c=US). • <i>e-mail-address</i>—The e-mail address of the new user to which OTPs and notices are to be sent.
Step 2	<pre>crypto ca server user-db allow user</pre> <p>Example:</p> <pre>hostname (config-ca-server)# crypto ca server user-db allow user6</pre>	<p>Provides user privileges to a newly added user.</p>
Step 3	<pre>crypto ca server user-db email-otp username</pre> <p>Example:</p> <pre>hostname (config-ca-server)# crypto ca server user-db email-otp exampleuser1</pre>	<p>Notifies a user in the local CA database to enroll and download a user certificate, which automatically e-mails the OTP to that user.</p> <p>Note When an administrator wants to notify a user through e-mail, the administrator must specify the e-mail address in the username field or in the e-mail field when adding that user.</p>

	Command	Purpose
Step 4	<pre>crypto ca server user-db show-otp</pre> <p>Example: <pre>hostname (config-ca-server)# crypto ca server user-db show-otp</pre></p>	Shows the issued OTP.
Step 5	<pre>otp expiration timeout</pre> <p>Example: <pre>hostname (config-ca-server)# otp expiration 24</pre></p>	<p>Sets the enrollment time limit in hours. The default expiration time is 72 hours. The otp expiration command defines the amount of time that the OTP is valid for user enrollment. This time period begins when the user is allowed to enroll.</p> <p>After a user enrolls successfully within the time limit and with the correct OTP, the local CA server creates a PKCS12 file, which includes a keypair for the user and a user certificate that is based on the public key from the keypair generated and the subject-name DN specified when the user is added. The PKCS12 file contents are protected by a passphrase, the OTP. The OTP can be handled manually, or the local CA can e-mail this file to the user to download after the administrator allows enrollment.</p> <p>The PKCS12 file is saved to temporary storage with the name, <i>username.p12</i>. With the PKCS12 file in storage, the user can return within the enrollment-retrieval time period to download the PKCS12 file as many times as needed. When the time period expires, the PKCS12 file is removed from storage automatically and is no longer available to download.</p> <p>Note If the enrollment period expires before the user retrieves the PKCS12 file that includes the user certificate, enrollment is not permitted.</p>

Renewing Users

To specify the timing of renewal notices, perform the following steps:

	Command	Purpose
Step 1	<pre>crypto ca server</pre> <p>Example: ciscoasa (config)# crypto ca server</p>	<p>Enters local ca server configuration mode. Allows you to configure and manage a local CA.</p>
Step 2	<pre>renewal-reminder time</pre> <p>Example: ciscoasa (config-ca-server)# renewal-reminder 7</p>	<p>Specifies the number of days (1-90) before the local CA certificate expires that an initial reminder to reenroll is sent to certificate owners. If a certificate expires, it becomes invalid.</p> <p>Renewal notices and the times they are e-mailed to users are variable, and can be configured by the administrator during local CA server configuration.</p> <p>Three reminders are sent. An e-mail is automatically sent to the certificate owner for each of the three reminders, provided an e-mail address is specified in the user database. If no e-mail address exists for the user, a syslog message alerts you of the renewal requirement.</p> <p>The ASA automatically grants certificate renewal privileges to any user who holds a valid certificate that is about to expire, as long as the user still exists in the user database. Therefore, if an administrator does not want to allow a user to renew automatically, the administrator must remove the user from the database before the renewal time period.</p>

Restoring Users

To restore a user and a previously revoked certificate that was issued by the local CA server, perform the following steps:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	crypto ca server unrevoke cert-serial-no Example: ciscoasa (config)# crypto ca server unrevoke 782ea09f	Restores a user and unrevokes a previously revoked certificate that was issued by the local CA server. The local CA maintains a current CRL with serial numbers of all revoked user certificates. This list is available to external devices and can be retrieved directly from the local CA if it is configured to do so with the cdp-url command and the publish-crl command. When you revoke (or unrevoke) any current certificate by certificate serial number, the CRL automatically reflects these changes.

Removing Users

To delete a user from the user database by username, perform the following steps:

	Command	Purpose
Step 1	crypto ca server Example: ciscoasa (config)# crypto ca server	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	crypto ca server user-db remove username Example: ciscoasa (config)# crypto ca server user-db remove user1	Removes a user from the user database and allows revocation of any valid certificates that were issued to that user.

Revoking Certificates

To revoke a user certificate, perform the following steps:

	Command	Purpose
Step 1	<pre>crypto ca server</pre> <p>Example: ciscoasa (config)# crypto ca server</p>	Enters local ca server configuration mode. Allows you to configure and manage a local CA.
Step 2	<pre>crypto ca server revoke cert-serial-no</pre> <p>Example: ciscoasa (config-ca-server)# crypto ca server revoke 782ea09f</p>	<p>Enters the certificate serial number in hexadecimal format. Marks the certificate as revoked in the certificate database on the local CA server and in the CRL, which is automatically reissued.</p> <p>Note The password is also required if the certificate for the ASA needs to be revoked, so make sure that you record it and store it in a safe place.</p>

Maintaining the Local CA Certificate Database

To maintain the local CA certificate database, make sure that you save the certificate database file, LOCAL-CA-SERVER.cdb, with the **write memory** command each time that a change to the database occurs. The local CA certificate database includes the following files:

- The LOCAL-CA-SERVER.p12 file is the archive of the local CA certificate and keypair that is generated when the local CA server is initially enabled.
- The LOCAL-CA-SERVER.crl file is the actual CRL.
- The LOCAL-CA-SERVER.ser file keeps track of the issued certificate serial numbers.

Rolling Over Local CA Certificates

Thirty days before the local CA certificate expires, a rollover replacement certificate is generated, and a syslog message informs the administrator that it is time for local CA rollover. The new local CA certificate must be imported onto all necessary devices before the current certificate expires. If the administrator does not respond by installing the rollover certificate as the new local CA certificate, validations may fail.

The local CA certificate rolls over automatically after expiration using the same keypair. The rollover certificate is available for export in base 64 format.

Examples

The following example shows a base 64 encoded local CA certificate:

```
MIIXlwIBAzCCF1EGCSqGSIb3DQEHAaCCF0IEghe+MIIXOjCCFzYGCsGSIb3DQEHBqCCFycwghc jAgEAMIIXHAYJKo
ZlIhvcNAQcBMBsGCiqGSIb3DQEAMwDQIQIjph4SxJJoyTgCAQGAghbw3v4bFy+GGG2dJnB4OLphsUM+IG3SDoiDwZG9
n1SvtMieoxd7Hxknxbum06JDrujWktHBIqkrm+td34q1NE1iGeP2YC94/NQ2z+4kS+uZzwcRh11KEZTS1E4L0fSaC3
uMTxJq2NUHYWmoc8pi4CIeLj3h7VVMY6qbx2AC8I+q57+QG5vG515Hi5imwtYfaWwPEdPQxaWZPrzoG1J8BFqdPa1j
BGhAzzuSmElm3j/2dQ3Atro1G9nIsRHgV39fcBgwz4fEabHG7/Vanb+fj81d5n1oiJjDYYbP86tVbZ2yOVZR6aKFVI
0b2AfCr6Pbwfc9U8Z/aF3BCyM2sN2xPjrXva94CaYrQyotZdAkSYA5KWSyEcgdqmuBeGDKOncTknfgy0XM+fG5rb3
qAXy1GkjyFI5Bm9Do6RUR0oG1DSrQrKeq/hj...
```

END OF CERTIFICATE

Archiving the Local CA Server Certificate and Keypair

To archive the local CA server certificate and keypair, enter the following command:

Command	Purpose
copy Example: hostname# copy LOCAL-CA-SERVER_0001.p12 tftp://10.1.1.22/user6/	Copies the local CA server certificate and keypair and all files from the ASA using either FTP or TFTP. Note Make sure that you back up all local CA files as often as possible.



Note

Monitoring Digital Certificates

To display certificate configuration and database information, enter one or more of the following commands:

Command	Purpose
show crypto ca server	Shows local CA configuration and status.
show crypto ca server cert-db	Shows user certificates issued by the local CA.
show crypto ca server certificate	Shows local CA certificates on the console in base 64 format and the rollover certificate when available, including the rollover certificate thumbprint for verification of the new certificate during import onto other devices.
show crypto ca server crt	Shows CRLs.
show crypto ca server user-db	Shows users and their status, which can be used with the following qualifiers to reduce the number of displayed records: <ul style="list-style-type: none"> • allowed. Shows only users currently allowed to enroll. • enrolled. Shows only users that are enrolled and hold a valid certificate • expired. Shows only users holding expired certificates. • on-hold. Lists only users without a certificate and not currently allowed to enroll.
show crypto ca server user-db allowed	Shows users who are eligible to enroll.
show crypto ca server user-db enrolled	Shows enrolled users with valid certificates.
show crypto ca server user-db expired	Shows users with expired certificates.
show crypto ca server user-db on-hold	Shows users without certificates who are not allowed to enroll.
show crypto key <i>name of key</i>	Shows key pairs that you have generated.
show running-config	Shows local CA certificate map rules.

Examples

The following example shows an RSA general-purpose key:

```
ciscoasa/contexta(config)# show crypto key mypubkey
Key pair was generated at: 16:39:47 central Feb 10 2010
Key name: <Default-RSA-Key>
Usage: General Purpose Key
Modulus Size (bits): 1024
Key Data:

30819f30 0d06092a 864886f7 0d010101 05000381 8d003081 89028181 00ea51b7
0781848f 78bccac2 4a1b5b8d 2f3e30b4 4cae9f86 f4485207 159108c9 f5e49103
9eeb0f5d 45fd1811 3b4aafce 292b3b64 b4124a6f 7a777b08 75b88df1 8092a9f8
5508e9e5 2c271245 7fd1c0c3 3aaf1e04 c7c4efa4 600f4c4a 6afe56ad c1d2c01c
e08407dd 45d9e36e 8cc0bfef 14f9e6ac eca141e4 276d7358 f7f50d13 79020301 0001
Key pair was generated at: 16:34:54 central Feb 10 2010
```

The following example shows the local CA CRL:

```
hostname (config)# show crypto ca server crl
Certificate Revocation List:
  Issuer: cn=xx5520-1-3-2007-1
  This Update: 13:32:53 UTC Jan 4 2010
  Next Update: 13:32:53 UTC Feb 3 2010
  Number of CRL entries: 2
  CRL size: 270 bytes
Revoked Certificates:
  Serial Number: 0x6f
  Revocation Date: 12:30:01 UTC Jan 4 2010
  Serial Number: 0x47
  Revocation Date: 13:32:48 UTC Jan 4 2010
```

The following example shows one user on-hold:

```
hostname (config)# show crypto ca server user-db on-hold
username: wilma101
email: <None>
dn: <None>
allowed: <not allowed>
notified: 0
hostname (config)#
```

The following example shows output of the **show running-config** command, in which local CA certificate map rules appear:

```
crypto ca certificate map 1
  issuer-name co asc
  subject-name attr ou eq Engineering
```

Feature History for Certificate Management

Table 40-1 lists each feature change and the platform release in which it was implemented.

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.
Certificate management	7.2(1)	We introduced the following commands: issuer-name <i>DN-string</i> , revocation-check crl none , revocation-check crl , revocation-check none . We deprecated the following commands: crl {required optional nocheck} .

Table 40-1 Feature History for Certificate Management (continued)

Feature Name	Platform Releases	Feature Information
Certificate management	8.0(2)	<p>We introduced the following commands:</p> <p>cdp-url, crypto ca server, crypto ca server crl issue, crypto ca server revoke <i>cert-serial-no</i>, crypto ca server unrevoke <i>cert-serial-no</i>, crypto ca server user-db add <i>user [dn dn] [email e-mail-address]</i>, crypto ca server user-db allow {<i>username</i> all-unenrolled all-certholders} [<i>display-otp</i>] [<i>email-otp</i>] [<i>replace-otp</i>], crypto ca server user-db email-otp {<i>username</i> all-unenrolled all-certholders}, crypto ca server user-db remove <i>username</i>, crypto ca server user-db show-otp {<i>username</i> all-certholders all-unenrolled}, crypto ca server user-db write, [no] database path <i>mount-name directory-path</i>, debug crypto ca server [<i>level</i>], lifetime {ca-certificate certificate crl} <i>time</i>, no shutdown, otp expiration <i>timeout</i>, renewal-reminder <i>time</i>, show crypto ca server, show crypto ca server cert-db [<i>user username</i> allowed enrolled expired on-hold] [<i>serial certificate-serial-number</i>], show crypto ca server certificate, show crypto ca server crl, show crypto ca server user-db [expired allowed on-hold enrolled], show crypto key <i>name of key</i>, show running-config, shutdown.</p>
SCEP proxy	8.4(1)	<p>We introduced this feature, which provides secure deployment of device certificates from third-party CAs.</p> <p>We introduced the following commands:</p> <p>crypto ikev2 enable outside client-services port <i>portnumber</i>, scep-enrollment enable, scep-forwarding-url value <i>URL</i>, secondary-pre-fill-username clientless hide use-common-password <i>password</i>, secondary-pre-fill-username ssl-client hide use-common-password <i>password</i>, secondary-username-from-certificate {use-entire-name use-script {<i>primary_attr</i> [<i>secondary_attr</i>]}} [no-certificate-fallback cisco-secure-desktop machine-unique-id].</p>



PART 8

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 41-1](#)
- [Configuring CLI Parameters, page 41-6](#)
- [Configuring ICMP Access, page 41-10](#)
- [Configuring Management Access Over a VPN Tunnel, page 41-13](#)
- [Configuring AAA for System Administrators, page 41-14](#)
- [Feature History for Management Access, page 41-37](#)



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 41-1](#)
- [Guidelines and Limitations, page 41-2](#)
- [Configuring Telnet Access, page 41-3](#)
- [Using a Telnet Client, page 41-3](#)
- [Configuring SSH Access, page 41-4](#)
- [Using an SSH Client, page 41-5](#)
- [Configuring HTTPS Access for ASDM, page 41-6](#)

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 41-13.
- 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 the **aaa authentication ssh console LOCAL** command; then define a local user by entering the **username** command. 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 13-1.
- 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 41-1.
- 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 Telnet Access

To identify the client IP addresses allowed to connect to the ASA using Telnet, perform the following steps.

Detailed Steps

	Command	Purpose
Step 1	<pre>telnet source_IP_address mask source_interface</pre> <p>Example:</p> <pre>ciscoasa(config)# telnet 192.168.1.2 255.255.255.255 inside</pre>	<p>For each address or subnet, identifies the IP addresses from which the ASA accepts connections.</p> <p>If there is only one interface, you can configure Telnet to access that interface as long as the interface has a security level of 100.</p>
Step 2	<pre>telnet timeout minutes</pre> <p>Example:</p> <pre>ciscoasa(config)# telnet timeout 30</pre>	<p>Sets the duration for how long a Telnet session can be idle before the ASA disconnects the session.</p> <p>Set the timeout from 1 to 1440 minutes. The default is 5 minutes. The default duration is too short in most cases and should be increased until all pre-production testing and troubleshooting have been completed.</p>

Examples

The following example shows how to let a host on the inside interface with an address of 192.168.1.2 access the ASA:

```
ciscoasa(config)# telnet 192.168.1.2 255.255.255.255 inside
```

The following example shows how to allow all users on the 192.168.3.0 network to access the ASA on the inside interface:

```
ciscoasa(config)# telnet 192.168.3.0 255.255.255.0 inside
```

Using a Telnet Client

To gain access to the ASA CLI using Telnet, enter the login password set by the **password** command. (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 13-1.

If you configure Telnet authentication (see the [“Configuring Authentication for CLI and ASDM Access” section on page 41-20](#)), then enter the username and password defined by the AAA server or local database.

Configuring SSH Access

To identify the client IP addresses and define a user allowed to connect to the ASA using SSH, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	crypto key generate rsa modulus <i>modulus_size</i> Example: ciscoasa(config)# crypto key generate rsa modulus 1024	Generates an RSA key pair, which is required for SSH. The modulus value (in bits) is 512, 768, 1024, or 2048. The larger the key modulus size you specify, the longer it takes to generate an RSA key pair. We recommend a value of 1024.
Step 2	write memory Example: ciscoasa(config)# write memory	Saves the RSA keys to persistent flash memory.
Step 3	aaa authentication ssh console LOCAL	Enables local authentication for SSH access. You can alternatively configure authentication using a AAA server. See the “Configuring Authentication for CLI and ASDM Access” section on page 41-20 for more information.
Step 4	username <i>username</i> password <i>password</i>	Creates a user in the local database that can be used for SSH access.
Step 5	ssh source_IP_address mask <i>source_interface</i> Example: ciscoasa(config)# ssh 192.168.3.0 255.255.255.0 inside	For each address or subnet, identifies the IP addresses from which the ASA accepts connections, and the interface on which you can SSH. Unlike Telnet, you can SSH on the lowest security level interface.
Step 6	ssh timeout <i>minutes</i> Example: ciscoasa(config)# ssh timeout 30	(Optional) Sets the duration for how long an SSH session can be idle before the ASA disconnects the session. Set the timeout from 1 to 60 minutes. The default is 5 minutes. The default duration is too short in most cases, and should be increased until all pre-production testing and troubleshooting have been completed.
Step 7	ssh version <i>version_number</i> Example: ciscoasa(config)# ssh version 2	(Optional) Limits access to SSH version 1 or 2. By default, SSH allows both versions 1 and 2.

Examples

The following example shows how to generate RSA keys and let a host on the inside interface with an address of 192.168.1.2 access the ASA:

```
ciscoasa(config)# crypto key generate rsa modulus 1024
ciscoasa(config)# write memory
ciscoasa(config)# aaa authentication ssh console LOCAL
WARNING: local database is empty! Use 'username' command to define local users.
ciscoasa(config)# username exampleuser1 password examplepassword1
ciscoasa(config)# ssh 192.168.1.2 255.255.255.255 inside
ciscoasa(config)# ssh timeout 30
```

The following example shows how to allow all users on the 192.168.3.0 network to access the ASA on the inside interface:

```
ciscoasa(config)# ssh 192.168.3.0 255.255.255.0 inside
```

Using an SSH Client

In the SSH client on your management host, enter the username and password that you configured in the [“Configuring SSH Access” section on page 41-4](#). 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 33-4](#).

Configuring HTTPS Access for ASDM

To use ASDM, you need to enable the HTTPS server, and allow HTTPS connections to the ASA. HTTPS access is enabled as part of the factory default configuration or when you use the **setup** command. This section describes how to manually configure ASDM access.

To configure HTTPS access for ASDM, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>http source_IP_address mask source_interface</pre> <p>Example: ciscoasa(config)# http 192.168.1.2 255.255.255.255 inside</p>	For each address or subnet, identifies the IP addresses from which the ASA accepts HTTPS connections.
Step 2	<pre>http server enable [port]</pre> <p>Example: ciscoasa(config)# http server enable 443</p>	<p>Enables the HTTPS server.</p> <p>By default, the <i>port</i> is 443. If you change the port number, be sure to include it in the ASDM access URL. For example, if you change the port number to 444, enter the following:</p> <p>https://10.1.1.1:444</p>

Examples

The following example shows how to enable the HTTPS server and let a host on the inside interface with an address of 192.168.1.2 access ASDM:

```
ciscoasa(config)# http server enable
ciscoasa(config)# http 192.168.1.2 255.255.255.255 inside
```

The following example shows how to allow all users on the 192.168.3.0 network to access ASDM on the inside interface:

```
ciscoasa(config)# http 192.168.3.0 255.255.255.0 inside
```

Configuring CLI Parameters

This section includes the following topics:

- [Licensing Requirements for CLI Parameters, page 41-7](#)
- [Guidelines and Limitations, page 41-7](#)
- [Configuring a Login Banner, page 41-7](#)
- [Customizing a CLI Prompt, page 41-8](#)
- [Changing the Console Timeout, page 41-9](#)

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

Command	Purpose
banner { exec login motd } <i>text</i> Example: <pre>ciscoasa(config)# banner motd Welcome to \$(hostname).</pre>	<p>Adds a banner to display at one of three times: when a user first connects (message-of-the-day (motd)), when a user logs in (login), and when a user accesses privileged EXEC mode (exec). When a user connects to the ASA, the message-of-the-day banner appears first, followed by the login banner and prompts. After the user successfully logs in to the ASA, the exec banner appears.</p> <p>To add more than one line, precede each line by the banner command.</p> <p>For the banner text:</p> <ul style="list-style-type: none"> • Spaces are allowed, but tabs cannot be entered using the CLI. • There are no limits for banner length 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.

Examples

The following example shows how to add a message-of-the-day banner:

```
ciscoasa(config)# banner motd Welcome to $(hostname).
ciscoasa(config)# banner motd Contact me at admin@example.com for any
ciscoasa(config)# banner motd issues.
```

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	<p>Displays the traffic-passing state of the unit. The following values appear for the state:</p> <ul style="list-style-type: none"> 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. <p>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.</p>

Detailed Steps

To customize the CLI prompt, enter the following commands:

Command	Purpose
<pre>prompt {[hostname] [context] [domain] [slot] [state] [priority] [cluster-unit]}</pre> <p>Example: ciscoasa(config)# firewall transparent</p>	Customizes the CLI prompt.

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

Command	Purpose
<pre>console timeout <i>number</i></pre> <p>Example: ciscoasa(config)# console timeout 0</p>	Specifies the idle time in minutes (0 through 60) after which the privileged session ends. The default timeout is 0, which means the session does not time out.

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 41-10](#)
- [Licensing Requirements for ICMP Access, page 41-10](#)
- [Guidelines and Limitations, page 41-11](#)
- [Default Settings, page 41-11](#)
- [Configuring ICMP Access, page 41-12](#)

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, enter one of the following commands:

Detailed Steps

Command	Purpose
(For IPv4) icmp { permit deny } { host <i>ip_address</i> <i>ip_address mask</i> any } [<i>icmp_type</i>] <i>interface_name</i> Example: ciscoasa(config)# icmp deny host 10.1.1.15 inside	Creates an IPv4 ICMP access rule. If you do not specify an <i>icmp_type</i> , all types are identified. You can enter the number or the name. To control ping, specify echo-reply (0) (ASA-to-host) or echo (8) (host-to-ASA). See the “ICMP Types” section on page 49-15 for a list of ICMP types.
(For IPv6) ipv6 icmp { permit deny } { <i>ipv6-prefix/prefix-length</i> any host <i>ipv6-address</i> } [<i>icmp-type</i>] <i>interface_name</i> Example: ciscoasa(config)# icmp permit host fe80::20d:88ff:feee:6a82 outside	Creates an IPv6 ICMP access rule. If you do not specify an <i>icmp_type</i> , all types are identified. You can enter the number or the name. To control ping, specify echo-reply (0) (ASA-to-host) or echo (8) (host-to-ASA). See the “ICMP Types” section on page 49-15 for a list of ICMP types.

Examples

The following example shows how to allow all hosts except the one at 10.1.1.15 to use ICMP to the inside interface:

```
hostname(config)# icmp deny host 10.1.1.15 inside
hostname(config)# icmp permit any inside
```

The following example shows how to allow the host at 10.1.1.15 to use only ping to the inside interface, enter the following command:

```
hostname(config)# icmp permit host 10.1.1.15 inside
```

The following example shows how to deny all ping requests and permit all packet-too-big messages (to support path MTU discovery) at the outside interface:

```
hostname(config)# ipv6 icmp deny any echo-reply outside
hostname(config)# ipv6 icmp permit any packet-too-big outside
```

The following example shows how to permit host 2000:0:0:4::2 or hosts on prefix 2001::/64 to ping the outside interface:

```
hostname(config)# ipv6 icmp permit host 2000:0:0:4::2 echo-reply outside
hostname(config)# ipv6 icmp permit 2001::/64 echo-reply outside
hostname(config)# ipv6 icmp permit any packet-too-big outside
```

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 41-13](#)
- [Guidelines and Limitations, page 41-2](#)
- [Configuring a Management Interface, page 41-14](#)

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

Command	Purpose
<code>management-access management_interface</code>	The <i>management_interface</i> specifies the name of the management interface that you want to access when entering the ASA from another interface.
Example: <code>ciscoasa(config)# management-access inside</code>	

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 41-14](#)
- [Licensing Requirements for AAA for System Administrators, page 41-18](#)
- [Prerequisites, page 41-18](#)
- [Guidelines and Limitations, page 41-19](#)
- [Default Settings, page 41-19](#)
- [Configuring Authentication for CLI and ASDM Access, page 41-20](#)
- [Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\), page 41-21](#)
- [Limiting User CLI and ASDM Access with Management Authorization, page 41-23](#)
- [Configuring a Password Policy for Local Database Users, page 41-24](#)
- [Configuring Command Authorization, page 41-27](#)
- [Configuring Management Access Accounting, page 41-33](#)
- [Viewing the Currently Logged-In User, page 41-34](#)
- [Setting a Management Session Quota, page 41-35](#)
- [Exchanging Keys in an SSH Session, page 41-35](#)
- [Recovering from a Lockout, page 41-36](#)

Information About AAA for System Administrators

This section describes AAA for system administrators and includes the following topics:

- [Information About Management Authentication, page 41-15](#)

- [Information About Command Authorization, page 41-16](#)

Information About Management Authentication

This section describes authentication for management access and includes the following topics:

- [Comparing CLI Access with and without Authentication, page 41-15](#)
- [Comparing ASDM Access with and without Authentication, page 41-15](#)
- [Authenticating Sessions from the Switch to the ASA Services Module, page 41-16](#)

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 (set with the **password** command). (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 (set by the **enable password** command). However, if you do not use enable authentication, after you enter the **enable** command, you are no longer logged in as a particular user. To maintain your username, use enable authentication.
- **Authentication**—If you configure enable authentication (see the [Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\), page 41-21](#)), 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. See the [“Authenticating Users with the login Command” section on page 41-22](#) for more information.

Comparing ASDM Access with and without Authentication

By default, you can log into ASDM with a blank username and the enable password set by the **enable password** command. 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 41-16](#)
- [About Preserving User Credentials, page 41-16](#)
- [Security Contexts and Command Authorization, page 41-17](#)

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 41-27). (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 the applicable AAA server-type chapter.
- Local Database—See the [“Adding a User Account to the Local Database”](#) section on page 33-4.

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 41-1.

Prerequisites for Local Command Authorization

- Configure **enable** authentication. (See the [“Configuring Authentication for CLI and ASDM Access”](#) section on page 41-20.)

enable authentication is essential for maintaining the username after the user accesses the **enable** command.

Alternatively, you can use the **login** command (which is the same as the **enable** command with authentication; for the local database only), which requires no configuration. We do not recommend this option because it is not as secure as **enable** authentication.

You can also use CLI authentication, but it is not required.

- See the following prerequisites for each user type:
 - Local database users—Configure each user in the local database at a privilege level from 0 to 15.
 - RADIUS users—Configure the user with Cisco VSA CVPN3000-Privilege-Level with a value between 0 and 15.

- LDAP users—Configure the user with a privilege level between 0 and 15, and then map the LDAP attribute to Cisco VSA CVPN3000-Privilege-Level according to the “[Configuring LDAP Attribute Maps](#)” section on page 36-5.

Prerequisites for TACACS+ Command Authorization

- Configure CLI authentication (see the “[Configuring Authentication for CLI and ASDM Access](#)” section on page 41-20).
- Configure **enable** authentication (see the “[Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\)](#)” section on page 41-21).

Prerequisites for Management Accounting

- Configure CLI authentication (see the “[Configuring Authentication for CLI and ASDM Access](#)” section on page 41-20).
- Configure **enable** authentication (see the “[Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\)](#)” section on page 41-21).

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 41-31.

Configuring Authentication for CLI and ASDM 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 41-1.
- For SSH access, you must configure SSH authentication; there is no default username.

Detailed Steps

	Command	Purpose
Step 1	<pre>aaa authentication {telnet ssh http serial} console {LOCAL server_group [LOCAL]}</pre> <p>Example:</p> <pre>ciscoasa(config)# aaa authentication ssh console radius_1 LOCAL ciscoasa(config)# aaa authentication http console radius_1 LOCAL ciscoasa(config)# aaa authentication serial console LOCAL</pre>	<p>Authenticates users for management access. The telnet keyword controls Telnet access. For the ASASM, this keyword 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 41-16.</p> <p>The ssh keyword controls SSH access.</p> <p>The http keyword controls ASDM access.</p> <p>The serial keyword controls console port access. For the ASASM, this keyword 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 41-16.</p> <p>HTTP management authentication does not support the SDI protocol for a AAA server group.</p> <p>If you use a AAA server group for authentication, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Specify the server group name followed by LOCAL (LOCAL is case sensitive). 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.</p> <p>You can alternatively use the local database as your primary method of authentication (with no fallback) by entering LOCAL alone.</p>
Step 2	<pre>http authentication-certificate interface</pre> <p>Example:</p> <pre>http authentication-certificate inside</pre>	<p>Requires a certificate from ASDM clients connecting over HTTP on the specified interface. This command can be used in addition to the aaa authentication command for ASDM.</p> <p>This command is only for ASDM access, use the command ssl certificate-authentication to require a certificate for all other SSL traffic, for example, cut-through proxy.</p>

Configuring Authentication to Access Privileged EXEC Mode (the enable Command)

You can configure the ASA to authenticate users with a AAA server or the local database when they enter the **enable** command. Alternatively, users are automatically authenticated with the local database when they enter the **login** command, which also accesses privileged EXEC mode depending on the user level in the local database.

This section includes the following topics:

- [Configuring Authentication for the enable Command, page 41-22](#)
- [Authenticating Users with the login Command, page 41-22](#)

Configuring Authentication for the enable Command

You can configure the ASA to authenticate users when they enter the **enable** command. See the [“Comparing CLI Access with and without Authentication”](#) section on page 41-15 for more information.

To authenticate users who enter the **enable** command, enter the following command.

Command	Purpose
<pre>aaa authentication enable console {LOCAL server_group [LOCAL]}</pre> <p>Example: <pre>ciscoasa(config)# aaa authentication enable console LOCAL</pre></p>	<p>Authenticates users who enter the enable command. The user is prompted for the username and password.</p> <p>If you use a AAA server group for authentication, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Specify the server group name followed by LOCAL (LOCAL is case sensitive). 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.</p> <p>You can alternatively use the local database as your primary method of authentication (with no fallback) by entering LOCAL alone.</p>

Authenticating Users with the login Command

From user EXEC mode, you can log in as any username in the local database using the **login** command.

This feature allows users to log in with their own username and password to access privileged EXEC mode, so you do not have to provide the system enable password to everyone. To allow users to access privileged EXEC mode (and all commands) when they log in, set the user privilege level to 2 (the default) through 15. If you configure local command authorization, then the user can only enter commands assigned to that privilege level or lower. See the [“Configuring Local Command Authorization”](#) section on page 41-27 for more information.



Caution

If you add users to the local database who can gain access to the CLI and whom you do not want to enter privileged EXEC mode, you should configure command authorization. Without command authorization, users can access privileged EXEC mode (and all commands) at the CLI using their own password if their privilege level is 2 or greater (2 is the default). Alternatively, you can use a AAA server for authentication, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged EXEC mode.

To log in as a user from the local database, enter the following command:

Command	Purpose
<pre>login</pre> <p>Example: <pre>ciscoasa# login</pre></p>	<p>Logs in as a user from the local database. The ASA prompts for your username and password. After you enter your password, the ASA places you in the privilege level that the local database specifies.</p>

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 configure the **aaa authentication serial console** command, then any user who authenticates can access the console port.

Detailed Steps

- Step 1** To enable management authorization for local, RADIUS, LDAP (mapped), and TACACS+ users, enter the following command:

```
ciscoasa(config)# aaa authorization exec {authentication-server | LOCAL}
```

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 41-27 for more information.

When the **authentication-server** option is configured, the same server is used for both authentication and authorization.

- 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 using the **ldap map-attributes** command. 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 **aaa authentication console** commands.
- Service-Type 7 (NAS prompt)—Allows access to the CLI when you configure the **aaa authentication {telnet | ssh} console** command, but denies ASDM configuration access if you configure the **aaa authentication http console** command. ASDM monitoring access is allowed. If you configure **enable** authentication with the **aaa authentication enable console** command, 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 **aaa authentication console** commands (excluding the **serial** keyword; 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.

The following example shows how to define an LDAP attribute map. In this example, the security policy specifies that users being authenticated through LDAP map the user record fields or parameters title and company to the IETF-RADIUS service-type and privilege-level, respectively.

```
ciscoasa(config)# ldap attribute-map admin-control
ciscoasa(config-ldap-attribute-map)# map-name title IETF-RADIUS-Service-Type
ciscoasa(config-ldap-attribute-map)# map-name company Privilege-Level
```

The following example applies an LDAP attribute map to an LDAP AAA server:

```
ciscoasa(config)# aaa-server ldap-server (dmz1) host 10.20.30.1
ciscoasa(config-aaa-server-host)# ldap-attribute-map admin-control
```

TACACS+ users

Authorization is requested with “service=shell,” and the server responds with PASS or FAIL.

- PASS, privilege level 1—Allows access to ASDM, with limited read-only access to the configuration and monitoring sections, and access for **show** commands that are privilege level 1 only.
- PASS, privilege level 2 and higher—Allows access to the CLI when you configure the **aaa authentication {telnet | ssh} console** command, but denies ASDM configuration access if you configure the **aaa authentication http console** command. ASDM monitoring access is allowed. If you configure **enable** authentication with the **aaa authentication enable console** command, 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 **aaa authentication console** commands(excluding the **serial** keyword; serial access is allowed).

Local users

Set the **service-type** command for a given username. By default, the **service-type** is **admin**, which allows full access to any services specified by the **aaa authentication console** command. For more information, see the [“Adding a User Account to the Local Database” section on page 33-4](#).

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 41-25](#)
- [Changing Your Password, page 41-27](#)

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 **username** command as well as the **change-password** command.

Prerequisites

- Configure CLI/ASDM authentication according to the “[Configuring Authentication for CLI and ASDM Access](#)” section on page 41-20. Be sure to specify the local database.
- Configure enable authentication according to the “[Configuring Authentication to Access Privileged EXEC Mode \(the enable Command\)](#)” section on page 41-21. Be sure to specify the local database.

Detailed Steps

	Command	Purpose
Step 1	<p>password-policy lifetime <i>days</i></p> <p>Example: ciscoasa(config)# password-policy lifetime 180</p>	<p>(Optional) Sets 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.</p> <p>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:</p> <ul style="list-style-type: none"> • Have another administrator change your password with the username command. • Log in to the physical console port to change your password.
Step 2	<p>password-policy minimum-changes <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-changes 2</p>	<p>(Optional) Sets 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.</p> <p>Character matching is position independent, meaning that new password characters are considered changed only if they do not appear anywhere in the current password.</p>
Step 3	<p>password-policy minimum-length <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-length 8</p>	<p>(Optional) Sets the minimum length of passwords. Valid values are between 3 and 64 characters. We recommend a minimum password length of 8 characters.</p>

	Command	Purpose
Step 4	<p>password-policy minimum-uppercase <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-uppercase 3</p>	(Optional) Sets the minimum number of upper case characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0, which means there is no minimum.
Step 5	<p>password-policy minimum-lowercase <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-lowercase 6</p>	(Optional) Sets the minimum number of lower case characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0, which means there is no minimum.
Step 6	<p>password-policy minimum-numeric <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-numeric 1</p>	(Optional) Sets the minimum number of numeric characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0, which means there is no minimum.
Step 7	<p>password-policy minimum-special <i>value</i></p> <p>Example: ciscoasa(config)# password-policy minimum-special 2</p>	(Optional) Sets the minimum number of special characters that passwords must have. Valid values are between 0 and 64 characters. Special characters include the following: !, @, #, \$, %, ^, &, *, '(and ')'. The default value is 0, which means there is no minimum.
Step 8	<p>password-policy authenticate <i>enable</i></p> <p>Example: ciscoasa(config)# password-policy authenticate enable</p>	<p>(Optional) Sets whether users must change their password using the change-password command, instead of letting users change their password with the username command. The default setting is disabled: a user can use either method to change their password.</p> <p>If you enable this feature, if you try to change your password with the username command, the following error message appears:</p> <pre>ERROR: Changing your own password is prohibited</pre> <p>You also cannot delete your own account with the clear configure username command. If you try, the following error message appears:</p> <pre>ERROR: You cannot delete all usernames because you are not allowed to delete yourself</pre>

Changing Your Password

If you configure a password lifetime in the password policy, you need to change your **username** password to a new one when the old password expires. This password change method is required if you enable password policy authentication (the **password-policy authenticate enable** command). If password policy authentication is not enabled, then you can use this method, or you can change your user account directly with the **username** command.

Detailed Steps

Command	Purpose
<pre>change-password [old-password old_password [new-password new_password]]</pre> <p>Example:</p> <pre>hostname# change-password old-password j0hncr1cht0n new-password a3rynsun</pre>	<p>Changes your username password. If you do not enter the old and new passwords in the command, the ASA prompts you for input.</p>

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 41-16.

This section includes the following topics:

- [Configuring Local Command Authorization, page 41-27](#)
- [Viewing Local Command Privilege Levels, page 41-31](#)
- [Configuring Commands on the TACACS+ Server, page 41-32](#)
- [Configuring TACACS+ Command Authorization, page 41-33](#)

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 33-4

- [“Supported Authentication Methods” section on page 34-1](#)
- [“Configuring LDAP Attribute Maps” section on page 36-5](#)

To configure local command authorization, perform the following steps:

Detailed Steps

Command	Purpose
<p>Step 1</p> <pre>privilege [show clear cmd] level level [mode {enable cmd}] command command</pre> <p>Example:</p> <pre>ciscoasa(config)# privilege show level 5 command filter</pre>	<p>Assigns a command to a privilege level.</p> <p>Repeat this command for each command that you want to reassign.</p> <p>The options in this command are the following:</p> <ul style="list-style-type: none"> • show clear cmd—These optional keywords let you 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. If you do not use one of these keywords, all forms of the command are affected. • level level—A level between 0 and 15. • mode {enable configure}—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: <ul style="list-style-type: none"> – enable—Specifies both user EXEC mode and privileged EXEC mode. – configure—Specifies configuration mode, accessed using the configure terminal command. • command command—The command you are configuring. You can only configure the privilege level of the <i>main</i> command. For example, you can configure the level of all aaa commands, but not the level of the aaa authentication command and the aaa authorization command separately.

	Command	Purpose
Step 2	<pre>aaa authorization exec authentication-server</pre> <p>Example:</p> <pre>ciscoasa(config)# aaa authorization exec authentication-server</pre>	<p>Supports administrative user privilege levels from RADIUS.</p> <p>Enforces user-specific access levels for users who authenticate for management access (see the aaa authentication console LOCAL command).</p> <p>Without this command, the ASA only supports privilege levels for local database users and defaults all other types of users to level 15.</p> <p>This command also enables management authorization for local, RADIUS, LDAP (mapped), and TACACS+ users.</p> <p>Use the aaa authorization exec LOCAL command to enable attributes to be taken from the local database. See the “Limiting User CLI and ASDM Access with Management Authorization” section on page 41-23 for information about configuring a user on a AAA server to accommodate management authorization.</p>
Step 3	<pre>aaa authorization command LOCAL</pre> <p>Example:</p> <pre>ciscoasa(config)# aaa authorization command LOCAL</pre>	<p>Enables the use of local command privilege levels, which can be checked with the privilege level of users in the local database, RADIUS server, or LDAP server (with mapped attributes).</p> <p>When you set command privilege levels, command authorization does not occur unless you configure command authorization with this command.</p>

Examples

The **filter** command has the following forms:

- **filter** (represented by the **configure** option)
- **show running-config filter**
- **clear configure filter**

You can set the privilege level separately for each form, or set the same privilege level for all forms by omitting this option. The following example shows how to set each form separately:

```
ciscoasa(config)# privilege show level 5 command filter
ciscoasa(config)# privilege clear level 10 command filter
ciscoasa(config)# privilege cmd level 10 command filter
```

Alternatively, the following example shows how to set all filter commands to the same level:

```
ciscoasa(config)# privilege level 5 command filter
```

The **show privilege** command separates the forms in the display.

The following example shows the use of the **mode** keyword. The **enable** command must be entered from user EXEC mode, while the **enable password** command, which is accessible in configuration mode, requires the highest privilege level:

```
ciscoasa(config)# privilege cmd level 0 mode enable command enable
ciscoasa(config)# privilege cmd level 15 mode cmd command enable
ciscoasa(config)# privilege show level 15 mode cmd command enable
```

The following example shows an additional command, the **configure** command, which uses the **mode** keyword:

```
ciscoasa(config)# privilege show level 5 mode cmd command configure
```

```
ciscoasa(config)# privilege clear level 15 mode cmd command configure
ciscoasa(config)# privilege cmd level 15 mode cmd command configure
ciscoasa(config)# privilege cmd level 15 mode enable command configure
```

**Note**

This last line is for the **configure terminal** command.

Viewing Local Command Privilege Levels

The following commands let you view privilege levels for commands.

Command	Purpose
show running-config all privilege all	Shows all commands.
show running-config privilege level <i>level</i>	Shows commands for a specific level. The <i>level</i> is an integer between 0 and 15.
show running-config privilege command <i>command</i>	Shows the level of a specific command.

Examples

For the **show running-config all privilege all** command, the ASA displays the current assignment of each CLI command to a privilege level. The following is sample output from this command:

```
ciscoasa(config)# show running-config all privilege all
privilege show level 15 command aaa
privilege clear level 15 command aaa
privilege configure level 15 command aaa
privilege show level 15 command aaa-server
privilege clear level 15 command aaa-server
privilege configure level 15 command aaa-server
privilege show level 15 command access-group
privilege clear level 15 command access-group
privilege configure level 15 command access-group
privilege show level 15 command access-list
privilege clear level 15 command access-list
privilege configure level 15 command access-list
privilege show level 15 command activation-key
privilege configure level 15 command activation-key
....
```

The following example shows the command assignments for privilege level 10:

```
ciscoasa(config)# show running-config privilege level 10
privilege show level 10 command aaa
```

The following example shows the command assignments for the **access-list** command:

```
ciscoasa(config)# show running-config privilege command access-list
privilege show level 15 command access-list
privilege clear level 15 command access-list
privilege configure level 15 command access-list
```

Configuring Commands on the TACACS+ Server

You can configure commands on a Cisco Secure Access Control Server (ACS) TACACS+ server as a shared profile component, for a group, or for individual users. For third-party TACACS+ servers, see your server documentation for more information about command authorization support.

See the following guidelines for configuring commands in Cisco Secure ACS Version 3.1; many of these guidelines also apply to third-party servers:

- The ASA sends the commands to be authorized as shell commands, so configure the commands on the TACACS+ server as shell commands.



Note Cisco Secure ACS might include a command type called “pix-shell.” Do not use this type for ASA command authorization.

- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

For example, to allow the **show running-configuration aaa-server** command, add **show running-configuration** to the command field, and type **permit aaa-server** in the arguments field.

- You can permit all arguments of a command that you do not explicitly deny by checking the **Permit Unmatched Args** check box.

For example, you can configure just the **show** command, 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.

- 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**.
- 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.

- 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.

- 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 41-36.

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 41-27.

To configure TACACS+ command authorization, enter the following command:

Detailed Steps

Command	Purpose
<pre>aaa authorization command tacacs+_server_group [LOCAL]</pre> <p>Example:</p> <pre>ciscoasa(config)# aaa authorization command group_1 LOCAL</pre>	<p>Performs command authorization using a TACACS+ server.</p> <p>You can configure the ASA to use the local database as a fallback method if the TACACS+ server is unavailable. To enable fallback, specify the server group name followed by LOCAL (LOCAL is case sensitive). We recommend that you use the same username and password in the local database as the TACACS+ 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 33-4) and command privilege levels (see the “Configuring Local Command Authorization” section on page 41-27).</p>

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

	Command	Purpose
Step 1	<pre>aaa accounting {serial telnet ssh enable} console server-tag</pre> <p>Example: ciscoasa(config)# aaa accounting telnet console group_1</p>	<p>Enables support for AAA accounting for administrative access. Valid server group protocols are RADIUS and TACACS+.</p>
Step 2	<pre>aaa accounting command [privilege level] server-tag</pre> <p>Example: ciscoasa(config)# aaa accounting command privilege 15 group_1</p>	<p>Enables command accounting. Only TACACS+ servers support command accounting.</p> <p>Where privilege level is the minimum privilege level and server-tag is the name of the TACACS+ server group to which the ASA should send command accounting messages.</p>

Viewing the Currently Logged-In User

To view the current logged-in user, enter the following command:

```
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 41-1 describes the **show curpriv** command output.

Table 41-1 *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: <ul style="list-style-type: none"> • 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, enter the following command:

Command	Purpose
<code>quota management-session number</code>	Sets the maximum number of simultaneous ASDM, SSH, and Telnet sessions that are allowed on the ASA. The no form of this command sets the quota value to 0, which means that there is no session limit.
Example: <code>hostname(config)# quota management-session 1000</code>	

Exchanging Keys in an SSH Session

The Diffie-Hellman (DH) key exchange provides a shared secret that cannot be determined by either party alone. The key exchange is combined with a signature and the host key to provide host authentication. This key-exchange method provides explicit server authentication.

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.

To exchange keys in an SSH session, enter the following command:

Command	Purpose
<code>ssh key-exchange group {dh-group1 dh-group14} sha-1</code>	Exchanges keys using either the DH Group 1 or DH Group 14 key-exchange method.
Example: <code>ciscoasa(config)# ssh key-exchange group dh-group14 sha-1</code> <code>ciscoasa# show running-config key-exchange</code> <code>ssh key-exchange dh-group14-sha1</code>	<p>The key-exchange keyword specifies that either the DH group 1 or DH group 14 key-exchange method will follow and should be used when exchanging keys.</p> <p>The group keyword indicates that either the DH group 1 key-exchange method or the DH group 14 key-exchange method will follow and should be used when exchanging keys.</p> <p>The dh-group1 keyword indicates that the DH group 1 key-exchange method will follow and should be used when exchanging keys. DH group 2 is called DH group 1 for legacy reasons.</p> <p>The dh-group14 keyword indicates that the DH group 14 key-exchange method will follow and should be used when exchanging keys.</p> <p>The sha-1 keyword indicates that the SHA-1 encryption algorithm should be used.</p> <p>Use the show running-config ssh key-exchange command to display the DH group key-exchange method currently being used.</p>

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 41-2](#) lists the common lockout conditions and how you might recover from them.

Table 41-2 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.	<ol style="list-style-type: none"> 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. 	<ol style="list-style-type: none"> 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.	<p>Fix the TACACS+ server user account.</p> <p>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.</p>	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.

Feature History for Management Access

Table 41-3 lists each feature change and the platform release in which it was implemented.

Table 41-3 Feature History for Management Access

Feature Name	Platform Releases	Feature Information
Management Access	7.0(1)	<p>We introduced this feature.</p> <p>We introduced the following commands:</p> <p>show running-config all privilege all, show running-config privilege level, show running-config privilege command, telnet, telnet timeout, ssh, ssh timeout, http, http server enable, asdm image disk, banner, console timeout, icmp, ipv6 icmp, management access, aaa authentication console, aaa authentication enable console, aaa authentication telnet ssh console, service-type, login, privilege, aaa authentication exec authentication-server, aaa authentication command LOCAL, aaa accounting serial telnet ssh enable console, show curpriv, aaa accounting command privilege.</p>
Increased SSH security; the SSH default username is no longer supported.	8.4(2)	<p>Starting in 8.4(2), you can no longer connect to the ASA using SSH with the <code>pix</code> or <code>asa</code> 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.</p>

Table 41-3 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)	<p>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.</p> <p>We introduced the following commands: change-password, password-policy lifetime, password-policy minimum changes, password-policy minimum-length, password-policy minimum-lowercase, password-policy minimum-uppercase, password-policy minimum-numeric, password-policy minimum-special, password-policy authenticate enable, clear configure password-policy, show running-config password-policy.</p>
Support for SSH public key authentication	8.4(4.1), 9.1(2)	<p>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).</p> <p>We introduced the following commands: ssh authentication.</p> <p><i>PKF key format support is only in 9.1(2) and later.</i></p>
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	8.4(4.1), 9.1(2)	<p>Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported.</p> <p>We introduced the following command: ssh key-exchange.</p>
Support for a maximum number of management sessions	8.4(4.1), 9.1(2)	<p>You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions.</p> <p>We introduced the following commands: quota management-session, show running-config quota management-session, show quota management-session.</p>
For the ASASM in multiple context mode, support for Telnet and virtual console authentication from the switch.	8.5(1)	<p>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.</p>
AES-CTR encryption for SSH	9.1(2)	<p>The SSH server implementation in the ASA now supports AES-CTR mode encryption.</p>

Table 41-3 Feature History for Management Access (continued)

Feature Name	Platform Releases	Feature Information
Improved SSH rekey interval	9.1(2)	An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic. We introduced the following command: show ssh sessions detail .
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 command: aaa authorization exec .



Managing Software and Configurations

This chapter describes how to manage the ASA software and configurations and includes the following sections:

- [Upgrading the Software, page 42-1](#)
- [Managing Files, page 42-12](#)
- [Configuring the Images and Startup Configuration to Use, page 42-21](#)
- [Using the ROM Monitor to Load an Image, page 42-22](#)
- [Backing Up Configurations or Other Files, page 42-25](#)
- [Downgrading Your Software, page 42-34](#)
- [Configuring Auto Update, page 42-35](#)

Upgrading the Software

This section describes how to upgrade to the latest version and includes the following topics:

- [Upgrade Path and Migrations, page 42-1](#)
- [Viewing Your Current Version, page 42-3](#)
- [Downloading the Software from Cisco.com, page 42-3](#)
- [Upgrading a Standalone Unit, page 42-3](#)
- [Upgrading a Failover Pair or ASA Cluster, page 42-5](#)

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.

- For a pre-9.1(2.8) version, to upgrade to 9.1(2.8) or 9.1(3) and later, you must be running one of the following versions:
 - 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:

Pre-9.1(2.8) 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 42-1](#) shows the supported scenarios for performing zero-downtime upgrades.

Table 42-1 Zero-Downtime Upgrade Support

Type of Upgrade	Support
Maintenance Release	<p>You can upgrade from any maintenance release to any other maintenance release within a minor release.</p> <p>For example, you can upgrade from 8.4(1) to 8.4(6) without first installing the maintenance releases in between.</p>

Table 42-1 Zero-Downtime Upgrade Support (continued)

Type of Upgrade	Support
Minor Release	<p>You can upgrade from a minor release to the next minor release. You cannot skip a minor release.</p> <p>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).</p> <hr/> <p> Note Zero-downtime upgrades are possible, even when feature configuration is migrated, for example, from 8.2 to 8.3.</p>
Major Release	<p>You can upgrade from the last minor release of the previous version to the next major release.</p> <p>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).</p> <hr/> <p> Note Zero-downtime upgrades are possible, even when feature configuration is migrated, for example, from 8.4 to 9.0.</p>

Viewing Your Current Version

Use the **show version** command to verify the software version of your ASA.

Downloading the Software from Cisco.com

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>

This procedure assumes you put the images on a TFTP server, although other server types are supported.

Upgrading a Standalone Unit

This section describes how to install the ASDM and operating system (OS) images.

Detailed Steps

This procedure uses TFTP. For FTP or HTTP, see the **copy** command.

	Command	Purpose
Step 1	<p>more system:running-config</p> <p>Example: hostname# more system:running-config</p>	<p>(If there is a configuration migration) The output shows the configuration on the terminal so that you can back up your configuration. Copy the output from this command, then paste the configuration in to a text file.</p> <p>For other methods of backing up, see the configuration guide.</p>
Step 2	<p>copy tftp://server[/path]/asa_image_name {disk0:/ disk1:/} [path/]asa_image_name</p> <p>Example: hostname# copy tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin</p>	<p>Copies the ASA software to the active unit flash memory. For other methods than TFTP, see the copy command.</p>
Step 3	<p>copy tftp://server[/path]/asdm_image_name {disk0:/ disk1:/} [path/]asdm_image_name</p> <p>Example: hostname# copy tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin</p>	<p>Copies the ASDM image to the active unit flash memory.</p>
Step 4	<p>configure terminal</p> <p>Example: hostname(config)# configure terminal</p>	<p>If you are not already in global configuration mode, accesses global configuration mode.</p>
Step 5	<p>show running-config boot system</p> <p>Example: hostname(config)# show running-config boot system boot system disk0:/cdisk.bin boot system disk0:/asa841-smp-k8.bin</p>	<p>Shows the current boot images configured (up to 4). The ASA uses the images in the order listed; if the first image is unavailable, the next image is used, and so on. You cannot insert a new image URL at the top of the list; to specify the new image to be first, you must remove any existing entries, and enter the image URLs in the order desired, according to Step 6 and Step 7.</p>
Step 6	<p>no boot system {disk0:/ disk1:/} [path/]asa_image_name</p> <p>Example: hostname(config)# no boot system disk0:/cdisk.bin hostname(config)# no boot system disk0:/asa841-smp-k8.bin</p>	<p>Removes any existing boot image configurations so that you can enter the new boot image as your first choice.</p>
Step 7	<p>boot system {disk0:/ disk1:/} [path/]asa_image_name</p> <p>Example: hostname(config)# boot system disk0://asa911-smp-k8.bin</p>	<p>Sets the ASA image to boot (the one you just uploaded).</p> <p>Repeat this command for any backup images that you want to use in case this image is unavailable. For example, you can re-enter the images that you previously removed in Step 6.</p>

	Command	Purpose
Step 8	<pre>asdm image {disk0:/ disk1:/} [path/] asdm_image_name</pre> <p>Example: hostname(config)# asdm image disk0:/asdm-711.bin </p>	Sets the ASDM image to use (the one you just uploaded). You can only configure one ASDM image to use, so you do not need to first remove the existing configuration.
Step 9	<pre>write memory</pre> <p>Example: hostname(config)# write memory </p>	Saves the new settings to the startup configuration.
Step 10	<pre>reload</pre> <p>Example: hostname# reload </p>	Reloads the ASA.

Upgrading a Failover Pair or ASA Cluster

- [Upgrading an Active/Standby Failover Pair, page 42-5](#)
- [Upgrading an Active/Active Failover Pair, page 42-8](#)
- [Upgrading an ASA Cluster, page 42-10](#)

Upgrading an Active/Standby Failover Pair

To upgrade the Active/Standby failover pair, perform the following steps.

Requirements

Perform these steps on the active unit.

Detailed Steps

	Command	Purpose
Step 1	<pre>more system:running-config</pre> <p>Example: active# more system:running-config </p>	(If there is a configuration migration) The output shows the configuration on the terminal so that you can back up your configuration. Copy the output from this command, then paste the configuration in to a text file. For other methods of backing up, see the configuration guide.
Step 2	<pre>copy tftp://server[/path]/asa_image_name {disk0:/ disk1:/} [path/] asa_image_name</pre> <p>Example: active# copy tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin </p>	Copies the ASA software to the active unit flash memory. For other methods than TFTP, see the copy command.

	Command	Purpose
Step 3	<pre>failover exec mate copy /noconfirm tftp://server[/path]/filename {disk0:/ disk1:/}[path/]filename</pre> <p>Example:</p> <pre>active# failover exec mate copy /noconfirm tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin</pre>	Copies the software to the standby unit; be sure to specify the same path as for the active unit.
Step 4	<pre>copy tftp://server[/path]/asdm_image_name {disk0:/ disk1:/}[path/]asdm_image_name</pre> <p>Example:</p> <pre>active# copy tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin</pre>	Copies the ASDM image to the active unit flash memory.
Step 5	<pre>failover exec mate copy /noconfirm tftp://server[/path]/asdm_image_name {disk0:/ disk1:/}[path/]asdm_image_name</pre> <p>Example:</p> <pre>active# failover exec mate copy /noconfirm tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin</pre>	Copies the ASDM image to the standby unit; be sure to specify the same path as for the active unit.
Step 6	<pre>configure terminal</pre> <p>Example:</p> <pre>active(config)# configure terminal</pre>	If you are not already in global configuration mode, accesses global configuration mode.
Step 7	<pre>show running-config boot system</pre> <p>Example:</p> <pre>hostname(config)# show running-config boot system boot system disk0:/cdisk.bin boot system disk0:/asa841-smp-k8.bin</pre>	Shows the current boot images configured (up to 4). The ASA uses the images in the order listed; if the first image is unavailable, the next image is used, and so on. You cannot insert a new image URL at the top of the list; to specify the new image to be first, you must remove any existing entries, and enter the image URLs in the order desired, according to Step 8 and Step 9 .
Step 8	<pre>no boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# no boot system disk0:/cdisk.bin hostname(config)# no boot system disk0:/asa841-smp-k8.bin</pre>	Removes any existing boot image configurations so that you can enter the new boot image as your first choice.
Step 9	<pre>boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# boot system disk0://asa911-smp-k8.bin</pre>	Sets the ASA image to boot (the one you just uploaded). Repeat this command for any backup images that you want to use in case this image is unavailable. For example, you can re-enter the images that you previously removed in Step 8 .

	Command	Purpose
Step 10	<pre>asdm image {disk0:/ disk1:/} [path/] asdm_image_name</pre> <p>Example: hostname(config)# asdm image disk0:/asdm-711.bin</p>	Sets the ASDM image to use (the one you just uploaded). You can only configure one ASDM image to use, so you do not need to first remove the existing configuration.
Step 11	<pre>write memory</pre> <p>Example: active(config)# write memory</p>	Saves the new settings to the startup configuration.
Step 12	<pre>failover reload-standby</pre> <p>Example: active# failover reload-standby</p>	<p>Reloads the standby unit to boot the new image.</p> <p>Wait for the standby unit to finish loading. Use the show failover command to verify that the standby unit is in the Standby Ready state.</p>
Step 13	<pre>no failover active</pre> <p>Example: active# no failover active</p>	Forces the active unit to fail over to the standby unit.
Step 14	<pre>reload</pre> <p>Example: active# reload</p>	Reloads the former active unit (now the new standby unit). If you want to restore this unit to be active after it reloads, enter the failover active command.

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. Also perform these steps on the *primary* unit.

Detailed Steps

	Command	Purpose
Step 1	<pre>more system:running-config</pre> <p>Example: primary# more system:running-config </p>	<p>(If there is a configuration migration) The output shows the configuration on the terminal so that you can back up your configuration. Copy the output from this command, then paste the configuration in to a text file.</p> <p>For other methods of backing up, see the configuration guide.</p>
Step 2	<pre>copy tftp://server[/path]/asa_image_name {disk0:/ disk1:/}[path]/asa_image_name</pre> <p>Example: primary# copy tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin </p>	<p>Copies the ASA software to the primary unit flash memory. For other methods than TFTP, see the copy command.</p>
Step 3	<pre>failover exec mate copy /noconfirm tftp://server[/path]/filename {disk0:/ disk1:/}[path]/filename</pre> <p>Example: primary# failover exec mate copy /noconfirm tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin </p>	<p>Copies the software to the secondary unit; be sure to specify the same path as for the primary unit.</p>
Step 4	<pre>copy tftp://server[/path]/asdm_image_name {disk0:/ disk1:/}[path]/asdm_image_name</pre> <p>Example: primary# copy tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin </p>	<p>Copies the ASDM image to the primary unit flash memory.</p>
Step 5	<pre>failover exec mate copy /noconfirm tftp://server[/path]/asdm_image_name {disk0:/ disk1:/}[path]/asdm_image_name</pre> <p>Example: primary# failover exec mate copy /noconfirm tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin </p>	<p>Copies the ASDM image to the secondary unit; be sure to specify the same path as for the active unit.</p>

	Command	Purpose
Step 6	<pre>failover active group 1 failover active group 2</pre> <p>Example:</p> <pre>primary# failover active group 1 primary# failover active group 2</pre>	Makes both failover groups active on the primary unit.
Step 7	<pre>configure terminal</pre> <p>Example:</p> <pre>primary(config)# configure terminal</pre>	If you are not already in global configuration mode, accesses global configuration mode.
Step 8	<pre>show running-config boot system</pre> <p>Example:</p> <pre>hostname(config)# show running-config boot system boot system disk0:/cdisk.bin boot system disk0:/asa841-smp-k8.bin</pre>	Shows the current boot images configured (up to 4). The ASA uses the images in the order listed; if the first image is unavailable, the next image is used, and so on. You cannot insert a new image URL at the top of the list; to specify the new image to be first, you must remove any existing entries, and enter the image URLs in the order desired, according to Step 9 and Step 10 .
Step 9	<pre>no boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# no boot system disk0:/cdisk.bin hostname(config)# no boot system disk0:/asa841-smp-k8.bin</pre>	Removes any existing boot image configurations so that you can enter the new boot image as your first choice.
Step 10	<pre>boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# boot system disk0://asa911-smp-k8.bin</pre>	Sets the ASA image to boot (the one you just uploaded). Repeat this command for any backup images that you want to use in case this image is unavailable. For example, you can re-enter the images that you previously removed in Step 9 .
Step 11	<pre>asdm image {disk0:/ disk1:/}[path/]asdm_image_name</pre> <p>Example:</p> <pre>hostname(config)# asdm image disk0:/asdm-711.bin</pre>	Sets the ASDM image to use (the one you just uploaded). You can only configure one ASDM image to use, so you do not need to first remove the existing configuration.
Step 12	<pre>write memory</pre> <p>Example:</p> <pre>primary(config)# write memory</pre>	Saves the new settings to the startup configuration.
Step 13	<pre>failover reload-standby</pre> <p>Example:</p> <pre>primary# failover reload-standby</pre>	Reloads the secondary unit to boot the new image. Wait for the secondary unit to finish loading. Use the show failover command to verify that both failover groups are in the Standby Ready state.

	Command	Purpose
Step 14	<pre>no failover active group 1 no failover active group 2</pre> <p>Example:</p> <pre>primary# no failover active group 1 primary# no failover active group 2</pre>	Forces both failover groups to become active on the secondary unit.
Step 15	<pre>reload</pre> <p>Example:</p> <pre>primary# reload</pre>	Reloads the primary unit. If the failover groups are configured with the preempt command, they automatically become active on their designated unit after the preempt delay has passed. If the failover groups are not configured with the preempt command, you can return them to active status on their designated units using the failover active group command.

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

	Command	Purpose
Step 1	<pre>more system:running-config</pre> <p>Example:</p> <pre>master# more system:running-config</pre>	(If there is a configuration migration) Backs up your configuration file. Copy the output from this command, then paste the configuration in to a text file. For other methods of backing up, see the configuration guide.
Step 2	<pre>cluster exec copy /noconfirm tftp://server[/path]/asa_image_name {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>master# cluster exec copy /noconfirm tftp://10.1.1.1/asa911-smp-k8.bin disk0:/asa911-smp-k8.bin</pre>	Copies the ASA software to all units in the cluster. For other methods than TFTP, see the copy command.
Step 3	<pre>cluster exec copy /noconfirm tftp://server[/path]/asdm_image_name {disk0:/ disk1:/}[path/]asdm_image_name</pre> <p>Example:</p> <pre>master# cluster exec copy /noconfirm tftp://10.1.1.1/asdm-711.bin disk0:/asdm-711.bin</pre>	Copies the ASDM image to all units in the cluster.
Step 4	<pre>configure terminal</pre> <p>Example:</p> <pre>master(config)# configure terminal</pre>	If you are not already in global configuration mode, accesses global configuration mode.

	Command	Purpose
Step 5	<pre>show running-config boot system</pre> <p>Example:</p> <pre>hostname(config)# show running-config boot system boot system disk0:/cdisk.bin boot system disk0:/asa841-smp-k8.bin</pre>	Shows the current boot images configured (up to 4). The ASA uses the images in the order listed; if the first image is unavailable, the next image is used, and so on. You cannot insert a new image URL at the top of the list; to specify the new image to be first, you must remove any existing entries, and enter the image URLs in the order desired, according to Step 6 and Step 7 .
Step 6	<pre>no boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# no boot system disk0:/cdisk.bin hostname(config)# no boot system disk0:/asa841-smp-k8.bin</pre>	Removes any existing boot image configurations so that you can enter the new boot image as your first choice.
Step 7	<pre>boot system {disk0:/ disk1:/}[path/]asa_image_name</pre> <p>Example:</p> <pre>hostname(config)# boot system disk0://asa911-smp-k8.bin</pre>	Sets the ASA image to boot (the one you just uploaded). Repeat this command for any backup images that you want to use in case this image is unavailable. For example, you can re-enter the images that you previously removed in Step 6 .
Step 8	<pre>asdm image {disk0:/ disk1:/}[path/]asdm_image_name</pre> <p>Example:</p> <pre>hostname(config)# asdm image disk0:/asdm-711.bin</pre>	Sets the ASDM image to use (the one you just uploaded). You can only configure one ASDM image to use, so you do not need to first remove the existing configuration.
Step 9	<pre>write memory</pre> <p>Example:</p> <pre>master(config)# write memory</pre>	Saves the new settings to the startup configuration.
Step 10	<pre>cluster exec unit slave-unit reload noconfirm</pre> <p>Example:</p> <pre>master# cluster exec unit unit2 reload noconfirm</pre>	<p>Reloads each slave unit when you repeat this command for each unit name. 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.</p> <p>To view member names, enter cluster exec unit ?, or enter the show cluster info command.</p>
Step 11	<pre>no enable</pre> <p>Example:</p> <pre>master(config)# no enable</pre>	<p>Disables clustering on the master unit. Wait for 5 minutes for a new master to be selected and traffic to stabilize.</p> <p>Do not enter write memory; when the master unit reloads, you want clustering to be enabled on it.</p>
Step 12	<pre>reload noconfirm</pre> <p>Example:</p> <pre>master# reload noconfirm</pre>	Reloads the master unit. A new election takes place for a new master unit. When the former master unit rejoins the cluster, it will be a slave.

Managing Files

- [Viewing Files in Flash Memory, page 42-12](#)
- [Deleting Files from Flash Memory, page 42-12](#)
- [Erasing the Flash File System, page 42-13](#)
- [Configuring File Access, page 42-13](#)
- [Copying a File to the ASA, page 42-17](#)
- [Copying a File to the Startup or Running Configuration, page 42-19](#)

Viewing Files in Flash Memory

You can view files in flash memory and see information about files as follows:

- To view files in flash memory, enter the following command:

```
ciscoasa# dir [disk0: | disk1:]
```

Enter **disk0:** for the internal flash memory. The **disk1:** keyword represents the external flash memory. The internal flash memory is the default.

For example:

```
hostname# dir
```

```
Directory of disk0:/
500  -rw-  4958208    22:56:20 Nov 29 2004  cdisk.bin
2513 -rw-   4634      19:32:48 Sep 17 2004  first-backup
2788 -rw-   21601    20:51:46 Nov 23 2004  backup.cfg
2927 -rw-  8670632    20:42:48 Dec 08 2004  asdmfile.bin
```

- To view extended information about a specific file, enter the following command:

```
hostname# show file information [path:] filename
```

The default path is the root directory of the internal flash memory (disk0:/).

For example:

```
hostname# show file information cdisk.bin
```

```
disk0:/cdisk.bin:
  type is image (XXX) []
  file size is 4976640 bytes version 7.0(1)
```

The file size listed is for example only.

Deleting Files from Flash Memory

You can remove files from flash memory that you no longer need. To delete a file from flash memory, enter the following command:

```
hostname# delete disk0: filename
```

By default, the file is deleted from the current working directory if you do not specify a path. You may use wildcards when deleting files. You are prompted with the filename to delete, and then you must confirm the deletion.

Erasing the Flash File System

To erase the flash file system, perform the following steps:

-
- Step 1** Connect to the ASA console port according to the instructions in the [“Accessing the ASA Services Module Command-Line Interface”](#) section on page 3-2 or the [“Accessing the Appliance Command-Line Interface”](#) section on page 3-1.
 - Step 2** Power off the ASA, then power it on.
 - Step 3** During startup, press the **Escape** key when you are prompted to enter ROMMON mode.
 - Step 4** Enter the **erase** command, which overwrites all files and erases the file system, including hidden system files.

```
rommon #1> erase [disk0: | disk1: | flash:]
```

Configuring File Access

- [Configuring the FTP Client Mode, page 42-13](#)
- [Configuring the ASA as a Secure Copy Server, page 42-14](#)
- [Customizing the ASA Secure Copy Client, page 42-14](#)
- [Configuring the ASA TFTP Client Path, page 42-16](#)

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

Command	Purpose
<code>ftp mode passive</code>	Sets the FTP mode to passive.
Example: <code>ciscoasa(config)# ftp mode passive</code>	

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 41-1](#).
- The ASA license must have the strong encryption (3DES/AES) license to support SSH Version 2 connections.

Detailed Steps

Command	Purpose
<code>ssh scopy enable</code>	Enables the SCP server.
Example: <code>ciscoasa(config)# ssh scopy enable</code>	

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 [“Copying a File to the ASA” section on page 42-17](#)). This section lets you customize the SCP client operation.

Prerequisites

For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the `changeto system` command.

Detailed Steps

	Command	Purpose
<p>Step 1</p> <pre>[no] ssh stricthostkeycheck</pre> <p>Example:</p> <pre>ciscoasa# ssh stricthostkeycheck ciscoasa# copy x scp://cisco@10.86.95.9/x The authenticity of host '10.86.95.9 (10.86.95.9)' can't be established. RSA key fingerprint is dc:2e:b3:e4:e1:b7:21:eb:24:e9:37:81:cf:bb: c3:2a. Are you sure you want to continue connecting (yes/no)? yes Warning: Permanently added '10.86.95.9' (RSA) to the list of known hosts. Source filename [x]? Address or name of remote host [10.86.95.9]? Destination username [cisco]? Destination password []? cisco123 Destination filename [x]?</pre>	<p>Enables or disables SSH host key checking. 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.</p>	
<p>Step 2</p> <pre>ssh pubkey-chain [no] server ip_address {key-string key_string exit key-hash {md5 sha256} fingerprint}</pre> <p>Example:</p> <pre>ciscoasa(config)# ssh pubkey-chain ciscoasa(config-ssh-pubkey-chain)# server 10.7.8.9 ciscoasa(config-ssh-pubkey-server)# key-string Enter the base 64 encoded RSA public key. End with the word "exit" on a line by itself ciscoasa(config-ssh-pubkey-server-string)# c1:b1:30:29:d7:b8:de:6c:97:77:10:d7:46:41: 63:87 ciscoasa(config-ssh-pubkey-server-string)# exit ciscoasa(config-ssh-pubkey-server)# show running-config ssh pubkey-chain ssh pubkey-chain server 10.7.8.9 key-hash sha256 f1:22:49:47:b6:76:74:b2:db:26:fb:13:65:d8: 99:19:e7:9e:24:46:59:be:13:7f:25:27:70:9b: 0e:d2:86:12</pre>	<p>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.</p> <p>For each server, you can specify the key-string (public key) or key-hash (hashed value) of the SSH host.</p> <p>The <i>key_string</i> 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 <code>.ssh/id_rsa.pub</code> file. After you submit the Base64 encoded public key, that key is then hashed via SHA-256.</p> <p>The key-hash {md5 sha256} fingerprint enters the already hashed key (using an MD5 or SHA-256 key); for example, a key that you copied from show command output.</p>	

Examples

The following example adds an already hashed host key for the server at 10.86.94.170:

```
ciscoasa(config)# ssh pubkey-chain
ciscoasa(config-ssh-pubkey-chain)# server 10.86.94.170
ciscoasa(config-ssh-pubkey-server)# key-hash sha256
65:d9:9d:fe:1a:bc:61:aa:64:9d:fc:ee:99:87:38:df:a8:8e:d9:e9:ff:42:de:e8:8d:2d:bf:a9:2b:85:2e:19
```

The following example adds a host string key for the server at 10.7.8.9:

```
ciscoasa(config)# ssh pubkey-chain
ciscoasa(config-ssh-pubkey-chain)# server 10.7.8.9
ciscoasa(config-ssh-pubkey-server)# key-string
Enter the base 64 encoded RSA public key.
End with the word "exit" on a line by itself
ciscoasa(config-ssh-pubkey-server-string)# c1:b1:30:29:d7:b8:de:6c:97:77:10:d7:46:41:63:87
ciscoasa(config-ssh-pubkey-server-string)# exit
```

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 “Copying a File to the ASA” section on page 42-17 and “Backing Up Configurations or Other Files” section on page 42-25). 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

Command	Purpose
<pre>tftp-server interface_name server_ip filename</pre> <p>Example:</p> <pre>ciscoasa(config)# tftp-server inside 10.1.4.7 files/config1.cfg ciscoasa(config)# copy tftp: test.cfg</pre> <p>Address or name of remote host [10.1.4.7]?</p> <p>Source filename [files/config1.cfg]?config2.cfg</p> <p>Destination filename [test.cfg]?</p> <p>Accessing tftp://10.1.4.7/files/config2.cfg;int=outside...</p>	<p>Pre-defines the TFTP server address and filename for use with configure net and copy commands. You can override the filename when you enter the command; for example, when you use the copy command, you can take advantage of the pre-defined TFTP server address but still enter any filename at the interactive prompts.</p> <p>For the copy command, enter tftp: to use the tftp-server value instead of tftp://url.</p>

Copying a File to the ASA

This section describes how to copy the application image, ASDM software, a configuration file, or any other file that needs to be downloaded to internal or external flash memory from a TFTP, FTP, SMB, HTTP, HTTPS, or SCP server.

Guidelines

- 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.
- You cannot have two files with the same name but with different letter case in the same directory in flash memory. For example, if you attempt to download the file, Config.cfg, to a location that contains the file, config.cfg, you receive the following error message:

```
%Error opening disk0:/Config.cfg (File exists).
```

- For information about installing the Cisco SSL VPN client, see the *Cisco AnyConnect VPN Client Administrator Guide*. For information about installing Cisco Secure Desktop on the ASA, see the *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administrators*.
- To configure the ASA to use a specific application image or ASDM image if you have more than one installed, or have installed them in external flash memory, see the [“Configuring the Images and Startup Configuration to Use”](#) section on page 42-21.
- For multiple context mode, you must be in the system execution space.

Detailed Steps

Command	Purpose
<pre>copy [/noconfirm] tftp://server[/path]/src_filename {disk0 disk1}:[/path/]dest_filename</pre> <p>Example: ciscoasa# copy tftp://10.1.1.67/files/context1.cfg disk0:/context1.cfg</p> <p>Address or name of remote host [10.1.1.67]? Source filename [files/context1.cfg]? Destination filename [context1.cfg]? Cryptochecksum: db8ba196 9ad189a8 7f5f501f 1bec469b !!!!!!!!!!!! 11143 bytes copied in 5.710 secs (2228 bytes/sec)</p>	Copies from a TFTP server.
<pre>copy [/noconfirm] ftp://[user[:password]@]server[/path]/src_filename {disk0 disk1}:[/path/]dest_filename</pre> <p>Example: ciscoasa# copy ftp://jcrichton:aeryn@10.1.1.67/files/context1.cfg disk0:/contexts/context1.cfg</p> <p>Address or name of remote host [10.1.1.67]? Source username [jcrichton]? Source password [aeryn]? Source filename [files/context1.cfg]? Destination filename [contexts/context1.cfg]? Cryptochecksum: db8ba196 9ad189a8 7f5f501f 1bec469b !!!!!!!!!!!! 11143 bytes copied in 5.710 secs (2228 bytes/sec)</p>	Copies from an FTP server.
<pre>copy [/noconfirm] http[s]://[user[:password]@]server[:port][/path]/src_filename {disk0 disk1}:[/path/]dest_filename</pre> <p>Example: ciscoasa# copy https://asun:john@10.1.1.67/files/moya.cfg disk0:/contexts/moya.cfg</p> <p>Address or name of remote host [10.1.1.67]? Source username [asun]? Source password [john]? Source filename [files/moya.cfg]? Destination filename [contexts/moya.cfg]? Cryptochecksum: db8ba196 9ad189a8 7f5f501f 1bec469b !!!!!!!!!!!! 11143 bytes copied in 5.710 secs (2228 bytes/sec)</p>	Copies from an HTTP(S) server.

Command	Purpose
<pre>copy [/noconfirm] smb://[user[:password]@]server[/path]/src_filename {disk0 disk1}:[/path/]dest_filename</pre> <p>Example:</p> <pre>ciscoasa# copy /noconfirm smb://chiana:dargo@10.1.1.67/test.xml disk0:/test.xml</pre> <pre>Cryptochecksum: db8ba196 9ad189a8 7f5f501f 1bec469b !!!!!!!!!!!! 11143 bytes copied in 5.710 secs (2228 bytes/sec)</pre>	Copies from an SMB server.
<pre>copy [/noconfirm] scp://[user[:password]@]server[/path]/src_filename[;int=interface_name] {disk0 disk1}:[/path/]dest_filename</pre> <p>Example:</p> <pre>ciscoasa# copy scp://pilot@10.86.94.170/test.cfg disk0:/test.cfg</pre> <pre>Address or name of remote host [10.86.94.170]? Source username [pilot]? Destination filename [test.cfg]? The authenticity of host '10.86.94.170 (10.86.94.170)' can't be established. RSA key fingerprint is <65:d9:9d:fe:1a:bc:61:aa:64:9d:fc:ee:99:87:38:df:a8:8e:d9:e9:ff:42:de:e8:8d: 2d:bf:a9:2b:85:2e:19> (SHA256) . Are you sure you want to continue connecting (yes/no)? yes Please use the following commands to add the hash key to the configuration: ssh pubkey-chain server 10.86.94.170 key-hash sha256 65:d9:9d:fe:1a:bc:61:aa:64:9d:fc:ee:99:87:38:df:a8:8e:d9:e9:ff:42:de:e8:8d:2 d:bf:a9:2b:85:2e:19 Password: <type in password> !!!!!!!! 6006 bytes copied in 8.160 secs (750 bytes/sec)</pre>	Copies from a SCP server. The int=interface option bypasses the route lookup and always uses the specified interface to reach the SCP server.

Copying a File to the Startup or Running Configuration

You can download a text file to the running or startup configuration from a TFTP, FTP, SMB, HTTP(S) or SCP server, or from the flash memory.

To configure the ASA to use a specific configuration as the startup configuration, see the [“Configuring the File to Boot as the Startup Configuration”](#) section on page 42-22.

Guidelines

When you copy a configuration to the running configuration, you merge the two configurations. 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.

Detailed Steps

To copy a file to the startup configuration or running configuration, enter one of the following commands for the appropriate download server:

Command	Purpose
<pre>copy [/noconfirm] tftp://server[/path]/src_filename {startup-config running-config}</pre> <p>Example: ciscoasa# copy tftp://10.1.1.67/files/old-running.cfg running-config</p>	Copies from a TFTP server.
<pre>copy [/noconfirm] ftp://[user[:password]@]server[/path]/src_filename {startup-config running-config}</pre> <p>Example: ciscoasa# copy ftp://jcrichon:aeryn@10.1.1.67/files/old-startup.cfg startup-config</p>	Copies from an FTP server.
<pre>copy [/noconfirm] http[s]://[user[:password]@]server[:port][/path]/src_filename {startup-config running-config}</pre> <p>Example: ciscoasa# copy https://asun:john@10.1.1.67/files/new-running.cfg running-config</p>	Copies from an HTTP(S) server.
<pre>copy [/noconfirm] smb://[user[:password]@]server[/path]/src_filename {startup-config running-config}</pre> <p>Example: ciscoasa# copy /noconfirm smb://chiana:dargo@10.1.1.67/new-running.cfg running-config</p>	Copies from an SMB server.
<pre>copy [/noconfirm] scp://[user[:password]@]server[/path]/src_filename[;int=interface_name] {startup-config running-config}</pre> <p>Example: ciscoasa# copy scp://pilot:moya@10.86.94.170/new-startup.cfg startup-config</p>	Copies from a SCP server. The ;int=interface option bypasses the route lookup and always uses the specified interface to reach the SCP server.

Examples

For example, to copy the configuration from a TFTP server, enter the following command:

```
ciscoasa# copy tftp://209.165.200.226/configs/startup.cfg startup-config
```


To copy the configuration from an FTP server, enter the following command:

```
ciscoasa# copy ftp://admin:letmein@209.165.200.227/configs/startup.cfg startup-config
```

To copy the configuration from an HTTP server, enter the following command:

```
ciscoasa# copy http://209.165.200.228/configs/startup.cfg startup-config
```

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 inserts the **asdm image** command into 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.

- [Configuring the ASA and ASDM Images to Use, page 42-21](#)
- [Configuring the File to Boot as the Startup Configuration, page 42-22](#)

Configuring the ASA and ASDM Images to Use

To configure the application image to boot, enter the following command:

```
ciscoasa(config)# boot system url
```

where *url* can be one of the following:

- `{disk0:/ | disk1:/}[path/]filename`
- `tftp://[user[:password]@]server[:port]/[path/]filename`



Note The TFTP option is not supported on all models.

You can enter up to four **boot system** command entries to specify different images to boot from in order; the ASA boots the first image it finds successfully. When you enter the **boot system** command, it adds an entry at the bottom of the list. To reorder the boot entries, you must remove all entries using the **clear configure boot system** command, and re-enter them in the order you desire. Only one **boot system tftp** command can be configured, and it must be the first one configured.



Note If the ASA is stuck in a cycle of constant booting, you can reboot the ASA into ROMMON mode. For more information about the ROMMON mode, see the [“Viewing Debugging Messages” section on page 43-1](#).

To configure the ASDM image to boot, enter the following command:

```
ciscoasa(config)# asdm image {disk0:/ | disk1:/}[path/]filename
```

Configuring the File to Boot as the Startup Configuration

By default, the ASA boots from a startup configuration that is a hidden file. You can alternatively set any configuration to be the startup configuration by entering the following command:

```
ciscoasa(config)# boot config {disk0:/ | disk1:/}[path/]filename
```

Using the ROM Monitor to Load an Image

- [Using ROM Monitor for the ASA 5500 Series, page 42-22](#)
- [Using the ROM Monitor for the ASASM, page 42-23](#)

Using ROM Monitor for the ASA 5500 Series

To load a software image to an ASA from the ROM monitor mode using TFTP, perform the following steps:

-
- Step 1** Connect to the ASA console port according to the instructions in the [“Accessing the Appliance Command-Line Interface”](#) section on page 3-1.
- Step 2** Power off the ASA, then power it on.
- Step 3** During startup, press the **Escape** key when you are prompted to enter ROMMON mode.
- Step 4** In ROMMON mode, define the interface settings to the ASA, including the IP address, TFTP server address, gateway address, software image file, and port, as follows:

```
rommon #1> ADDRESS=10.132.44.177
rommon #2> SERVER=10.129.0.30
rommon #3> GATEWAY=10.132.44.1
rommon #4> IMAGE=f1/asa800-232-k8.bin
rommon #5> PORT=Ethernet0/0
Ethernet0/0
Link is UP
MAC Address: 0012.d949.15b8
```



Note Be sure that the connection to the network already exists.

- Step 5** To validate your settings, enter the **set** command.

```
rommon #6> set
ROMMON Variable Settings:
  ADDRESS=10.132.44.177
  SERVER=10.129.0.30
  GATEWAY=10.132.44.1
  PORT=Ethernet0/0
  VLAN=untagged
  IMAGE=f1/asa840-232-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20
```

- Step 6** Ping the TFTP server by entering the **ping server** command.

```
rommon #7> ping server
Sending 20, 100-byte ICMP Echoes to server 10.129.0.30, timeout is 4 seconds:

Success rate is 100 percent (20/20)
```

Step 7 Load the software image by entering the **tftp** command.

```
rommon #8> tftp
ROMMON Variable Settings:
  ADDRESS=10.132.44.177
  SERVER=10.129.0.30
  GATEWAY=10.132.44.1
  PORT=Ethernet0/0
  VLAN=untagged
  IMAGE=f1/asa840-232-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20

tftp f1/asa840-232-k8.bin@10.129.0.30 via 10.132.44.1

Received 14450688 bytes

Launching TFTP Image...
Cisco ASA Security Appliance admin loader (3.0) #0: Mon Mar 5 16:00:07 MST 2011

Loading...N
After the software image is successfully loaded, the ASA automatically exits ROMMON mode.
```

Step 8 To verify that the correct software image has been loaded into the ASA, check the version in the ASA by entering the following command:

```
hostname# show version
```

Using the ROM Monitor for the ASASM

To load a software image to an ASASM from the ROM monitor mode using TFTP, perform the following steps:

- Step 1** Connect to the ASA console port according to the instructions in the [“Accessing the ASA Services Module Command-Line Interface”](#) section on page 3-2.
- Step 2** Make sure that you reload the ASASM image.
- Step 3** During startup, press the **Escape** key when you are prompted to enter ROMMON mode.
- Step 4** In ROMMON mode, define the interface settings to the ASASM, including the IP address, TFTP server address, gateway address, software image file, port, and VLAN, as follows:

```
rommon #1> ADDRESS=172.16.145.149
rommon #2> SERVER=172.16.171.125
rommon #3> GATEWAY=172.16.145.129
rommon #4> IMAGE=f1/asa851-smp-k8.bin
rommon #5> PORT=Data0
rommon #6> VLAN=1
Data0
Link is UP
MAC Address: 0012.d949.15b8
```



Note Be sure that the connection to the network already exists.

Step 5 To validate your settings, enter the **set** command.

```
rommon #7> set
ROMMON Variable Settings:
  ADDRESS=172.16.145.149
  SERVER=172.16.171.125
  GATEWAY=172.16.145.129
  PORT=Data0
  VLAN=1
  IMAGE=f1/asa851-smp-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=2
  RETRY=20
```

Step 6 Ping the TFTP server by entering the **ping server** command.

```
rommon #8> ping server
Sending 20, 100-byte ICMP Echoes to server 172.16.171.125, timeout is 2 seconds:

Success rate is 100 percent (20/20)
```

Step 7 Load the software image by entering the **tftp** command.

```
rommon #9> tftp
Clearing EOBC receive queue ...
cmostime_set = 1
ROMMON Variable Settings:
  ADDRESS=172.16.145.149
  SERVER=172.16.171.125
  GATEWAY=172.16.145.129
  PORT=Data0
  VLAN=1
  IMAGE=f1/asa851-smp-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=2
  RETRY=20

tftp f1/asa851-smp-k8.bin@172.16.171.125 via 172.16.145.129
Starting download. Press ESC to abort.
```

After the software image is successfully loaded, the ASASM automatically exits ROMMON mode.



Note You must download the image to the system flash separately after ROMMON boot is complete; booting the module into ROMMON mode does not preserve the system image across reloads.

Step 8 To verify that the correct software image has been loaded into the ASASM, check the version by entering the following command:

```
hostname# show version
```

Backing Up Configurations or Other Files

- [Backing up the Single Mode Configuration or Multiple Mode System Configuration, page 42-25](#)
- [Backing Up a Context Configuration or Other File in Flash Memory, page 42-26](#)
- [Backing Up a Context Configuration within a Context, page 42-27](#)
- [Copying the Configuration from the Terminal Display, page 42-27](#)
- [Backing Up Additional Files Using the Export and Import Commands, page 42-27](#)
- [Using a Script to Back Up and Restore Files, page 42-28](#)

Backing up the Single Mode Configuration or Multiple Mode System Configuration

In single context mode or from the system configuration in multiple mode, you can copy the startup configuration or running configuration to an external server or to the local flash memory.

Detailed Steps

Command	Purpose
<pre>copy [/noconfirm] {startup-config running-config} tftp://server[/path]/dst_filename</pre> <p>Example: ciscoasa# copy running-config tftp://10.1.1.67/files/new-running.cfg</p>	Copies to a TFTP server.
<pre>copy [/noconfirm] {startup-config running-config} ftp://[user[:password]@]server[/path]/dst_filename</pre> <p>Example: ciscoasa# copy startup-config ftp://jcrichton:aeryn@10.1.1.67/files/new-startup.cfg</p>	Copies to an FTP server.
<pre>copy [/noconfirm] {startup-config running-config} smb://[user[:password]@]server[/path]/dst_filename</pre> <p>Example: ciscoasa# copy /noconfirm running-config smb://chiana:dargo@10.1.1.67/new-running.cfg</p>	Copies to an SMB server.

Command	Purpose
<pre>copy [/noconfirm] {startup-config running-config} scp://[user[:password]@]server[/path]/dst_filename[;int=interface_name]</pre> <p>Example:</p> <pre>ciscoasa# copy startup-config scp://pilot:moya@10.86.94.170/new-startup.cfg</pre>	Copies to a SCP server. The int=interface option bypasses the route lookup and always uses the specified interface to reach the SCP server.
<pre>copy [/noconfirm] {startup-config running-config} {disk0 disk1}:[path/]dst_filename</pre> <p>Example:</p> <pre>ciscoasa# copy /noconfirm running-config disk0:/new-running.cfg</pre>	Copies to the local flash memory. Be sure that the destination directory exists. If it does not exist, first create the directory using the mkdir command.

Backing Up a Context Configuration or Other File in Flash Memory

Copy context configurations or other files that are on the local flash memory by entering one of the following commands in the system execution space.

Detailed Steps

Command	Purpose
<pre>copy [/noconfirm] {disk0 disk1}:[path/]src_filename tftp://server[/path]/dst_filename</pre> <p>Example:</p> <pre>ciscoasa# copy disk0:/asa-os.bin tftp://10.1.1.67/files/asa-os.bin</pre>	Copies from flash to a TFTP server.
<pre>copy [/noconfirm] {disk0 disk1}:[path/]src_filename ftp://[user[:password]@]server[/path]/dst_filename</pre> <p>Example:</p> <pre>ciscoasa# copy disk0:/asa-os.bin ftp://jcrichon:aeryn@10.1.1.67/files/asa-os.bin</pre>	Copies from flash to an FTP server.
<pre>copy [/noconfirm] {disk0 disk1}:[path/]src_filename smb://[user[:password]@]server[/path]/dst_filename</pre> <p>Example:</p> <pre>ciscoasa# copy /noconfirm copy disk0:/asdm.bin smb://chiana:dargo@10.1.1.67/asdm.bin</pre>	Copies from flash to an SMB server.

Command	Purpose
<pre>copy [/noconfirm] {disk0 disk1}:[path/]src_filename scp://[user[:password]@]server[/path]/dst_filename[;int=interface_name]</pre> <p>Example:</p> <pre>ciscoasa# copy disk0:/context1.cfg scp://pilot:moya@10.86.94.170/context1.cfg</pre>	<p>Copies from flash to SCP server. The ;int=interface option bypasses the route lookup and always uses the specified interface to reach the SCP server.</p>
<pre>copy [/noconfirm] {disk0 disk1}:[path/]src_filename {disk0 disk1}:[path/]dst_filename</pre> <p>Example:</p> <pre>ciscoasa# copy /noconfirm disk1:/file1.cfg disk0:/file1.cfgnew-running.cfg</pre>	<p>Copies from flash to the local flash memory. Be sure that the destination directory exists. If it does not exist, first create the directory using the mkdir command.</p>

Backing Up a Context Configuration within a Context

In multiple context mode, from within a context, you can perform the following backups:

- To copy the running configuration to the startup configuration server (connected to the admin context), enter the following command:

```
ciscoasa/contexta# copy running-config startup-config
```

- To copy the running configuration to a TFTP server connected to the context network, enter the following command:

```
ciscoasa/contexta# copy running-config tftp:/server[/path]/filename
```

Copying the Configuration from the Terminal Display

To print the configuration to the terminal, enter the following command:

```
ciscoasa# show running-config
```

Copy the output from this command, and then paste the configuration into a text file.

Backing Up Additional Files Using the Export and Import Commands

Additional files essential to your configuration might include the following:

- Files that you import using the **import webvpn** command. Currently, these files include customizations, URL lists, web content, plug-ins, and language translations.
- DAP policies (dap.xml).
- CSD configurations (data.xml).
- Digital keys and certificates.
- Local CA user database and certificate status files.

The CLI lets you back up and restore individual elements of your configuration using the **export** and **import** commands.

To back up these files, for example, those files that you imported with the **import webvpn** command or certificates, perform the following steps:

Step 1 Run the applicable **show** command(s) as follows:

```
ciscoasa # show import webvpn plug-in
ica
rdp
ssh, telnet
vnc
```

Step 2 Run the **export** command for the file that you want to back up (in this example, the rdp file):

```
ciscoasa # export webvpn plug-in protocol rdp tftp://tftpserver/backupfilename
```

Using a Script to Back Up and Restore Files

You can use a script to back up and restore the configuration files on your ASA, including all extensions that you import via the **import webvpn** CLI, the CSD configuration XML files, and the DAP configuration XML file. For security reasons, we do not recommend that you perform automated backups of digital keys and certificates or the local CA key.

This section provides instructions for doing so and includes a sample script that you can use as is or modify as your environment requires. The sample script is specific to a Linux system. To use it for a Microsoft Windows system, you need to modify it using the logic of the sample.



Note

The existing CLI lets you back up and restore individual files using the **copy**, **export**, and **import** commands. It does not, however, have a facility that lets you back up all ASA configuration files in one operation. Running the script facilitates the use of multiple CLIs.

This section includes the following topics:

- [Prerequisites, page 42-28](#)
- [Running the Script, page 42-29](#)
- [Sample Script, page 42-29](#)

Prerequisites

To use a script to back up and restore an ASA configuration, first perform the following tasks:

- Install Perl with an Expect module.
- Install an SSH client that can reach the ASA.
- Install a TFTP server to send files from the ASA to the backup site.

Another option is to use a commercially available tool. You can put the logic of this script into such a tool.

Running the Script

To run a backup-and-restore script, perform the following steps:

-
- Step 1** Download or cut-and-paste the script file to any location on your system.
 - Step 2** At the command line, enter **Perl** *scriptname*, where *scriptname* is the name of the script file.
 - Step 3** Press **Enter**.
 - Step 4** The system prompts you for values for each option. Alternatively, you can enter values for the options when you enter the **Perl** *scriptname* command before you press **Enter**. Either way, the script requires that you enter a value for each option.
 - Step 5** The script starts running, printing out the commands that it issues, which provides you with a record of the CLIs. You can use these CLIs for a later restore, which is particularly useful if you want to restore only one or two files.
-

Sample Script

```
#!/usr/bin/perl
#Function: Backup/restore configuration/extensions to/from a TFTP server.
#Description: The objective of this script is to show how to back up
configurations/extensions before the backup/restore command is developed.
# It currently backs up the running configuration, all extensions imported via "import
webvpn" command, the CSD configuration XML file, and the DAP configuration XML file.
#Requirements: Perl with Expect, SSH to the ASA, and a TFTP server.
#Usage: backupasa -option option_value
#       -h: ASA hostname or IP address
#       -u: User name to log in via SSH
#       -w: Password to log in via SSH
#       -e: The Enable password on the security appliance
#       -p: Global configuration mode prompt
#       -s: Host name or IP address of the TFTP server to store the configurations
#       -r: Restore with an argument that specifies the file name. This file is produced
during backup.
#If you don't enter an option, the script will prompt for it prior to backup.
#
#Make sure that you can SSH to the ASA.

use Expect;
use Getopt::Std;

#global variables
%options=();
$restore = 0; #does backup by default
$restore_file = '';
$asa = '';
$storage = '';
$user = '';
$password = '';
$enable = '';
$prompt = '';
$date = `date +%F`;
chop($date);
my $exp = new Expect();

getopts("h:u:p:w:e:s:r:", \%options);
do process_options();
```

```

do login($exp);
do enable($exp);
if ($restore) {
    do restore($exp,$restore_file);
}
else {
    $restore_file = "$prompt-restore-$date.cli";
    open(OUT,">$restore_file") or die "Can't open $restore_file\n";
    do running_config($exp);
    do lang_trans($exp);
    do customization($exp);
    do plugin($exp);
    do url_list($exp);
    do webcontent($exp);
    do dap($exp);
    do csd($exp);
    close(OUT);
}
do finish($exp);

sub enable {
    $obj = shift;
    $obj->send("enable\n");
    unless ($obj->expect(15, 'Password:')) {
        print "timed out waiting for Password:\n";
    }
    $obj->send("$enable\n");
    unless ($obj->expect(15, "$prompt#")) {
        print "timed out waiting for $prompt#\n";
    }
}

sub lang_trans {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("show import webvpn translation-table\n");
    $obj->expect(15, "$prompt#" );
    $output = $obj->before();
    @items = split(/\n+/, $output);

    for (@items) {
        s/^\s+//;
        s/\s+$//;
        next if /show import/ or /Translation Tables/;
        next unless (/^\.+s+.$/);
        ($lang, $transtable) = split(/\s+/, $_);
        $cli = "export webvpn translation-table $transtable language $lang
$storage/$prompt-$date-$transtable-$lang.po";
        $ocli = $cli;
        $ocli =~ s/^export/import/;
        print "$cli\n";
        print OUT "$ocli\n";
        $obj->send("$cli\n");
        $obj->expect(15, "$prompt#" );
    }
}

sub running_config {
    $obj = shift;
    $obj->clear_accum();
    $cli = "copy /noconfirm running-config $storage/$prompt-$date.cfg";
    print "$cli\n";
    $obj->send("$cli\n");
}

```

```

$obj->expect(15, "$prompt#" );
}

sub customization {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("show import webvpn customization\n");
    $obj->expect(15, "$prompt#" );
    $output = $obj->before();
    @items = split(/\n+/, $output);

    for (@items) {
        chop;
        next if /^Template/ or /show import/ or /^\s*$/;
        $cli = "export webvpn customization $_ $storage/$prompt-$date-cust-$_.xml";
        $ocli = $cli;
        $ocli =~ s/^export/import/;
        print "$cli\n";
        print OUT "$ocli\n";
        $obj->send("$cli\n");
        $obj->expect(15, "$prompt#" );
    }
}

sub plugin {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("show import webvpn plug-in\n");
    $obj->expect(15, "$prompt#" );
    $output = $obj->before();
    @items = split(/\n+/, $output);

    for (@items) {
        chop;
        next if /^Template/ or /show import/ or /^\s*$/;
        $cli = "export webvpn plug-in protocol $_ $storage/$prompt-$date-plugin-$_.jar";
        $ocli = $cli;
        $ocli =~ s/^export/import/;
        print "$cli\n";
        print OUT "$ocli\n";
        $obj->send("$cli\n");
        $obj->expect(15, "$prompt#" );
    }
}

sub url_list {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("show import webvpn url-list\n");
    $obj->expect(15, "$prompt#" );
    $output = $obj->before();
    @items = split(/\n+/, $output);

    for (@items) {
        chop;
        next if /^Template/ or /show import/ or /^\s*$/ or /No bookmarks/;
        $cli="export webvpn url-list $_ $storage/$prompt-$date-urllist-$_.xml";
        $ocli = $cli;
        $ocli =~ s/^export/import/;
        print "$cli\n";
        print OUT "$ocli\n";
        $obj->send("$cli\n");
        $obj->expect(15, "$prompt#" );
    }
}

```

```

    }
}

sub dap {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("dir dap.xml\n");
    $obj->expect(15, "$prompt#" );

    $output = $obj->before();
    return 0 if($output =~ /Error/);

    $cli="copy /noconfirm dap.xml $storage/$prompt-$date-dap.xml";
    $ocli="copy /noconfirm $storage/$prompt-$date-dap.xml disk0:/dap.xml";
    print "$cli\n";
    print OUT "$ocli\n";
    $obj->send("$cli\n");
    $obj->expect(15, "$prompt#" );
}

sub csd {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("dir sdesktop\n");
    $obj->expect(15, "$prompt#" );

    $output = $obj->before();
    return 0 if($output =~ /Error/);

    $cli="copy /noconfirm sdesktop/data.xml $storage/$prompt-$date-data.xml";
    $ocli="copy /noconfirm $storage/$prompt-$date-data.xml disk0:/sdesktop/data.xml";
    print "$cli\n";
    print OUT "$ocli\n";
    $obj->send("$cli\n");
    $obj->expect(15, "$prompt#" );
}

sub webcontent {
    $obj = shift;
    $obj->clear_accum();
    $obj->send("show import webvpn webcontent\n");
    $obj->expect(15, "$prompt#" );
    $output = $obj->before();
    @items = split(/\n+/, $output);

    for (@items) {
        s/^\s+//;
        s/\s+$//;
        next if /show import/ or /No custom/;
        next unless (/^\.+\.s+.$/);
        ($url, $type) = split(/\s+/, $_);
        $turl = $url;
        $turl =~ s/\/\+//;
        $turl =~ s/\/+\/-//;
        $cli = "export webvpn webcontent $url $storage/$prompt-$date-$turl";
        $ocli = $cli;
        $ocli =~ s/^export/import/;
        print "$cli\n";
        print OUT "$ocli\n";
        $obj->send("$cli\n");
        $obj->expect(15, "$prompt#" );
    }
}

```

```

sub login {
    $obj = shift;
    $obj->raw_pty(1);
    $obj->log_stdout(0); #turn off console logging.
    $obj->spawn("/usr/bin/ssh $user@$asa") or die "can't spawn ssh\n";
    unless ($obj->expect(15, "password:" )) {
        die "timeout waiting for password:\n";
    }

    $obj->send("$password\n");

    unless ($obj->expect(15, "$prompt>" )) {
        die "timeout waiting for $prompt>\n";
    }
}

sub finish {
    $obj = shift;
    $obj->hard_close();
    print "\n\n";
}

sub restore {
    $obj = shift;
    my $file = shift;
    my $output;
    open(IN,$file) or die "can't open $file\n";
    while (<IN>) {
        $obj->send("$_");
        $obj->expect(15, "$prompt#" );
        $output = $obj->before();
        print "$output\n";
    }
    close(IN);
}

sub process_options {
    if (defined($options{s})) {
        $tstr= $options{s};
        $storage = "tftp://$tstr";
    }
    else {
        print "Enter TFTP host name or IP address:";
        chop($tstr=<>);
        $storage = "tftp://$tstr";
    }
    if (defined($options{h})) {
        $asa = $options{h};
    }
    else {
        print "Enter ASA host name or IP address:";
        chop($asa=<>);
    }

    if (defined ($options{u})) {
        $user= $options{u};
    }
    else {
        print "Enter user name:";
        chop($user=<>);
    }

    if (defined ($options{w})) {

```

```

    $password= $options{w};
}
else {
    print "Enter password:";
    chop($password=<>);
}
if (defined ($options{p})) {
    $prompt= $options{p};
}
else {
    print "Enter ASA prompt:";
    chop($prompt=<>);
}
if (defined ($options{e})) {
    $enable = $options{e};
}
else {
    print "Enter enable password:";
    chop($enable=<>);
}

if (defined ($options{r})) {
    $restore = 1;
    $restore_file = $options{r};
}
}

```

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 42-34](#)
- [Performing the Downgrade, page 42-35](#)

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

To downgrade from Version 8.3, perform the following steps:

Detailed Steps

Step 1 Enter the following command:

```
ciscoasa(config)# downgrade [/noconfirm] old_image_url old_config_url [activation-key
old_key]
```

Where the **/noconfirm** option downgrades without prompting. The *image_url* is the path to the old image on disk0, disk1, tftp, ftp, or smb. The *old_config_url* is the path to the saved, premigration configuration (by default, this configuration was saved on disk0). If you need to revert to a pre-8.3 activation key, you can enter the old activation key.

This command 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 action 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**).

For example:

```
ciscoasa(config)# downgrade /noconfirm disk0:/asa821-k8.bin disk0:/8_2_1_0_startup_cfg.sav
```

Configuring Auto Update

This section includes the following topics:

- [Information About Auto Update, page 42-36](#)
- [Guidelines and Limitations, page 42-39](#)
- [Configuring Communication with an Auto Update Server, page 42-39](#)
- [Configuring Client Updates as an Auto Update Server, page 42-41](#)

- [Viewing Auto Update Status, page 42-42](#)

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 42-36](#)
- [Auto Update Benefits, page 42-36](#)
- [Auto Update Server Support in Failover Configurations, page 42-36](#)

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:


```

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 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 ASA as an Auto Update client, perform the following steps:

- Step 1** To specify the URL of the Auto Update Server, enter the following command:

```
ciscoasa(config)# auto-update server url [source interface] [verify-certificate]
```

where *url* has the following syntax:

```
http[s]://[user:password@]server_ip[:port]/pathname
```

SSL is used when **https** is specified. The *user* and *password* arguments of the URL are used for basic authentication when logging in to the server. If you use the **write terminal**, **show configuration** or **show tech-support** commands to view the configuration, the user and password are replaced with ‘*****’.

The default port is 80 for HTTP and 443 for HTTPS.

The **source interface** keyword and argument specify which interface to use when sending requests to the Auto Update Server. If you specify the same interface specified by the **management-access** command, the Auto Update requests travel over the same IPsec VPN tunnel used for management access.

The **verify-certificate** keyword verifies the certificate returned by the Auto Update Server.

- Step 2** (Optional) To identify the device ID to send when communicating with the Auto Update Server, enter the following command:

```
ciscoasa(config)# auto-update device-id {hardware-serial | hostname | ipaddress [if-name] | mac-address [if-name] | string text}
```

The identifier used is determined by specifying one of the following parameters:

- The *hardware-serial* argument specifies the ASA serial number.
- The *hostname* argument specifies the ASA hostname.
- The **ipaddress** keyword specifies the IP address of the specified interface. If the interface name is not specified, it uses the IP address of the interface used to communicate with the Auto Update Server.
- The **mac-address** keyword specifies the MAC address of the specified interface. If the interface name is not specified, it uses the MAC address of the interface used to communicate with the Auto Update Server.
- The **string** keyword specifies the specified text identifier, which cannot include white space or the characters ‘, “, , >, & and ?.

Step 3 (Optional) To specify how often to poll the Auto Update Server for configuration or image updates, enter the following command:

```
ciscoasa(config)# auto-update poll-period poll-period [retry-count [retry-period]]
```

The *poll-period* argument specifies how often (in minutes) to check for an update. The default is 720 minutes (12 hours).

The *retry-count* argument specifies how many times to try reconnecting to the server if the first attempt fails. The default is zero.

The *retry-period* argument specifies how long to wait (in minutes) between retries. The default is five minutes.

Step 4 (Optional) To schedule a specific time for the ASA to poll the Auto Update Server, enter the following command:

```
ciscoasa(config)# auto-update poll-at days-of-the-week time [randomize minutes]  
[retry_count [retry_period]]
```

The *days-of-the-week* argument is any single day or combination of days: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday. Other possible values are daily (Monday through Sunday), weekdays (Monday through Friday), and weekends (Saturday and Sunday).

The *time* argument specifies the time in the format HH:MM at which to start the poll. For example, 8:00 is 8:00 a.m. and 20:00 is 8:00 p.m.

The **randomize** *minutes* keyword and argument specify the period to randomize the poll time following the specified start time. The range is from 1 to 1439 minutes.

The *retry_count* argument specifies how many times to try reconnecting to the Auto Update Server if the first attempt fails. The default is zero.

The *retry_period* argument specifies how long to wait between connection attempts. The default is five minutes. The range is from 1 to 35791 minutes.

Step 5 (Optional) If the Auto Update Server has not been contacted for a certain period of time, entering the following command causes it to stop passing traffic:

```
ciscoasa(config)# auto-update timeout period
```

The *period* argument specifies the timeout period in minutes between 1 and 35791. The default is to never time out (zero minutes). To restore the default, enter the **no** form of this command.

Use the **auto-update timeout** command to be sure that the ASA has the most recent image and configuration. This condition is reported with system log message 201008.

In the following example, an ASA is configured to poll an Auto Update Server with the IP address 209.165.200.224, at port number 1742, from the outside interface, with certificate verification.

The ASA is also configured to use the hostname as the device ID and to poll an Auto Update Server every Friday and Saturday night at a random time between 10:00 p.m. and 11:00 p.m. On a failed polling attempt, the ASA will try to reconnect to the Auto Update Server ten times, and will wait three minutes between attempts at reconnecting, as shown in the following example:

```
ciscoasa(config)# auto-update server
https://jcrichton:farscape@209.165.200.224:1742/management source outside
verify-certificate
ciscoasa (config)# auto-update device-id hostname
hostname (config)# auto-update poll-at Friday Saturday 22:00 randomize 60 2 10
```

Configuring Client Updates as an Auto Update Server

Entering the **client-update** command enables updates for ASAs configured as Auto Update clients and lets you specify the type of software component (ASDM or boot image), the type or family of ASA, revision numbers to which the update applies, and a URL or IP address from which to obtain the update.

To configure the ASA as an Auto Update Server, perform the following steps:

Step 1 To enable client update, enter the following command:

```
ciscoasa(config)# client-update enable
```

Step 2 Configure the following parameters for the **client-update** command that you want to apply to the ASAs:

```
client-update {component {asdm | image} | device-id dev_string |
family family_name | type type} url url-string rev-nums rev-nums}
```

The **component** {**asdm** | **image**} parameter specifies the software component, either ASDM or the boot image of the ASA.

The **device-id** *dev_string* parameter specifies a unique string that the Auto Update client uses to identify itself. The maximum length is 63 characters.

The **family** *family_name* parameter specifies the family name that the Auto Update client uses to identify itself. It can be asa, pix, or a text string with a maximum length of seven characters.

The **rev-nums** *rev-nums* parameter specifies the software or firmware images for this client. Enter up to four, in any order, separated by commas.

The **type** *type* parameter specifies the type of clients to notify of a client update. Because this command is also used to update Windows clients, the list of clients includes several Windows operating systems. The ASAs in the list may include the following:

- asa5505: Cisco 5505 ASA
- asa5510: Cisco 5510 ASA
- asa5520: Cisco 5520 ASA
- asa5540: Cisco 5540 ASA

The **url** *url-string* parameter specifies the URL for the software/firmware image. This URL must point to a file appropriate for this client. For all Auto Update clients, you must use the protocol “http://” or “https://” as the prefix for the URL.

Configure the parameters for the client update that you want to apply to all ASAs of a particular type. That is, specify the type of ASA and the URL or IP address from which to get the updated image. In addition, you must specify a revision number. If the revision number of the remote ASA matches one of the specified revision numbers, there is no need to update the client, and the update is ignored.

To configure a client update for Cisco 5520 ASAs, enter the following command:

```
ciscoasa(config)# client-update type asa5520 component asdm url  
http://192.168.1.114/aus/asdm601.bin rev-nums 8.0(1)
```

Viewing Auto Update Status

To view the Auto Update status, enter the following command:

```
ciscoasa(config)# show auto-update
```

The following is sample output from the **show auto-update** command:

```
ciscoasa(config)# show auto-update  
  
Server: https://*****@209.165.200.224:1742/management.cgi?1276  
Certificate will be verified  
Poll period: 720 minutes, retry count: 2, retry period: 5 minutes  
Timeout: none  
Device ID: host name [corporate]  
Next poll in 4.93 minutes  
Last poll: 11:36:46 PST Tue Nov 13 2004  
Last PDM update: 23:36:46 PST Tue Nov 12 2004
```

Feature History for Software and Configurations

Table 42-2 lists each feature change and the platform release in which it was implemented.

Table 42-2 Feature History for Software and Configurations

Feature Name	Platform Releases	Feature Information
Secure Copy client	9.1(5)	<p>The ASA now supports the Secure Copy (SCP) client to transfer files to and from a SCP server.</p> <p>We introduced the following commands: ssh pubkey-chain, server (ssh pubkey-chain), key-string, key-hash, ssh stricthostkeycheck.</p> <p>We modified the following command: copy scp.</p>



Troubleshooting

This chapter describes how to troubleshoot the ASA and includes the following sections:

- [Viewing Debugging Messages, page 43-1](#)
- [Capturing Packets, page 43-2](#)
- [Viewing the Crash Dump, page 43-6](#)
- [Viewing the CoreDump, page 43-6](#)

Viewing Debugging Messages

Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC. Moreover, it is best to use **debug** commands during periods

of less network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system use. To enable debugging messages, see the **debug** commands in the command reference.

Capturing Packets

Capturing packets may be useful when troubleshooting connectivity problems or monitoring suspicious activity. We recommend that you contact Cisco TAC if you want to use the packet capture feature.

To capture packets, enter the following command:

Command	Purpose
<pre>[cluster exec] capture capture_name [type {asp-drop drop-code raw-data isakmp [ikev1 ikev2] tls-proxy lcp webvpn user user_name [form-only]}} [access-list acl_name] [buffer buf-size] [ethernet-type type] {interface {if-name asa_dataplane cluster}} [packet-length bytes] [circular-buffer] [headers-only] [match protocol {host source_ip source_ip mask any} [operator port] {host dest_ip dest_ip mask any} [operator port]] [real-time [dump] [detail] [trace]] [reinject-hide] [trace [detail] [trace-count number]]]</pre>	<p>Enables packet capture capabilities for packet sniffing and network fault isolation. For the complete syntax description, see the command reference or the CLI help (help capture). Not all options can be specified in one command. See the CLI help for allowed combinations.</p> <p>Use the same <i>capture_name</i> on multiple capture statements to capture multiple types of traffic.</p> <p>The type asp-drop keyword captures packets dropped by the accelerated security path. In a cluster, dropped forwarded data packets from one unit to another are also captured. In multiple context mode, when this option is issued in the system, all context dropped data packets are captured.</p> <p>The buffer keyword defines the buffer size used to store the packet. When the byte buffer is full, packet capture stops. When used in a cluster, this is the per-unit size, not the sum of all units.</p> <p>The circular-buffer keyword overwrites the buffer, starting from the beginning, when the buffer is full.</p> <p>The interface keyword sets the name of the interface on which to use packet capture. You must configure an interface for any packets to be captured.</p> <p>To capture packets on the dataplane, use the asa_dataplane keyword. To filter packets captured on the ASA CX backplane, use the asa_dataplane option and follow these guidelines. In single mode, the backplane control packets bypass the access list and are captured. In multiple context mode, only control packets are captured in the system context. Data packets are captured in the user context. The access-list and match options are only available in the user context.</p> <p>To capture the traffic on the cluster control link, use the cluster keyword. If you configure type lcp, specify the physical interface ID instead of the nameif name.</p> <p>The match keyword captures matching the protocol and source and destination IP addresses and optional ports. You can use this keyword up to three times in one command. The <i>operator</i> can be as follows:</p> <ul style="list-style-type: none"> • lt—less than • gt—greater than • eq—equal to <p>The type raw-data keywords capture inbound and outbound packets. This setting is the default.</p> <p>The real-time keyword displays the captured packets continuously in real-time. To terminate real-time packet capture, enter Ctrl + c. To permanently remove the capture, use the no form of this command. This option applies only to raw-data and asp-drop captures. This option is not supported when you use the cluster exec capture command.</p> <p>The reinject-hide keyword specifies that no reinjected packets will be captured and applies only in a clustering environment.</p> <p>Note If ACL optimization is configured, you cannot use the access-list command in capture. You can only use the access-group command. An error appears if you try to use the access-list command in this case.</p>
<p>Example: ciscoasa# capture captest interface inside</p>	

Capturing Packets in a Clustering Environment

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the master unit using the **cluster exec capture** command, which is then automatically enabled on all of the slave units in the cluster. The **cluster exec** keywords are the new keywords that you place in front of the **capture** command to enable cluster-wide capture.

The “cluster” interface name is the default name for the cluster control link and is not configurable. You specify “cluster “ as the interface name to capture the traffic on the cluster control link interface. There are two types of packets on the cluster control link: control plane packets and data plane packets, which both include forwarded data traffic and cluster LU messages. The TTL field in the IP address header is encoded to differentiate between these two types of packets. When forwarded data packets are captured, their clustering trailers are included in the capture file for debugging purposes.

In multiple context mode, although the cluster interface belongs to the system context, you can see the interface, so you can configure captures on the cluster link in user contexts. In the system context, both control plane and data plane packets are available. The data plane captures LU packets and forwarded data packets that belong only to the system context. In user contexts, control plane packets are not visible. Only forwarded data packets that belong to a specified user context and LU packets are captured. For security purposes, each context can only see the packets that belong to it.

Guidelines and Limitations

This section includes the guidelines and limitation for this feature.

Most of the limitations are the result of the distributed nature of the ASA architecture and the hardware accelerators that are being used in the ASA.

- You can only capture IP traffic; you cannot capture non-IP packets such as ARPs.
- For cluster control link capture in multiple context mode, only the packet that is associated with the context sent in the cluster control link is captured.
- In multicontext mode, the **copy capture** command is available only in the system space. The syntax is as follows:

```
copy /pcap capture:Context-name/in-cap tftp:
```

Where *in-cap* is the capture configured in the context *context-name*

- The **cluster exec capture realtime** command is not supported. The following error message appears:

```
Error: Real-time capture can not be run in cluster exec mode.
```
- For a shared VLAN, the following guidelines apply:
 - You can only configure one capture for the VLAN; if you configure a capture in multiple contexts on the shared VLAN, then only the last capture that was configured is used.
 - If you remove the last-configured (active) capture, no captures become active, even if you have previously configured a capture in another context; you must remove the capture and add it again to make it active.
 - All traffic that enters the interface to which the capture is attached is captured, including traffic to other contexts on the shared VLAN.
 - Therefore, if you enable a capture in Context A for a VLAN that is also used by Context B, both Context A and Context B ingress traffic are captured.

- For egress traffic, only the traffic of the context with the active capture is captured. The only exception is when you do not enable the ICMP inspection (therefore the ICMP traffic does not have a session in the accelerated path). In this case, both ingress and egress ICMP traffic for all contexts on the shared VLAN is captured.
- Configuring a capture typically involves configuring an ACL that matches the traffic that needs to be captured. After an ACL that matches the traffic pattern is configured, then you need to define a capture and associate this ACL to the capture, along with the interface on which the capture needs to be configured.

After you have performed a cluster-wide capture, to copy the same cluster-wide capture file to a TFTP server, enter the following command on the master unit:

```
ciscoasa (cfg-cluster)# cluster exec copy /pcap capture: cap_name
tftp://location/path/filename.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 filename_A.pcap, filename_B.pcap, and so on. In this example, A and B are cluster unit names. A different destination name is generated if you add the unit name at the end of the filename.

To enable cluster-wide capture on a specified interface, you can add the **cluster exec** keywords in front of each of the commands shown in the examples. These **capture** commands can only be replicated from the master unit to the slave units. However, you can still configure a capture on the specified interface for the local unit using any of these **capture** commands.

Examples

The following example shows how to create a cluster-wide LACP capture:

```
ciscoasa (config)# cluster exec capture lacp type lacp interface gigabitEthernet0/0
```

The following example shows how to create a capture for control path packets in the clustering link:

```
ciscoasa (config)# capture cp interface cluster match udp any eq 49495 any
ciscoasa (config)# capture cp interface cluster match udp any any eq 49495
```

The following example shows how to create a capture for data path packets in the clustering link:

```
ciscoasa (config)# access-list ccl extended permit udp any any eq 4193
ciscoasa (config)# access-list ccl extended permit udp any eq 4193 any
ciscoasa (config)# capture dp interface cluster access-list ccl
```

The following example shows how to capture data path traffic through the cluster:

```
ciscoasa (config)# capture abc interface inside match tcp host 1.1.1.1 host 2.2.2.2 eq www
ciscoasa (config)# capture abc interface inside match udp host 1.1.1.1 any
ciscoasa (config)# capture abc interface inside access-list xxx
```

The following example shows how to capture logical update messages for flows that match the real source to the real destination, and capture packets forwarded over CCL that match the real source to the real destination:

```
ciscoasa (config)# access-list dp permit ip real_src real_dst
```

The following example shows how to capture a certain type of data plane message, such as icmp echo request/response, that is forwarded from one ASA to another ASA using the **match** keyword or the ACL for the message type:

```
ciscoasa (config)# capture capture_name interface cluster access-list match icmp any any
```

The following example shows how to create a capture by using ACL 103 on a cluster control link:

```
ciscoasa (config)# access-list 103 permit ip A B
ciscoasa (config)# capture example1 interface cluster access-list 103
```

In the previous example, if A and B are IP addresses for the CCL interface, only the packets that are sent between these two units are captured.

If A and B are IP addresses for through-device traffic, then the following is true:

- Forwarded packets are captured as usual, provided the source and destination IP addresses are matched with the ACL.
- The data path logic update message is captured provided it is for the flow between A and B or for an ACL (for example, access-list 103). The capture matches the five-tuple of the embedded flow.
- Although the source and destination addresses in the UDP packet are CCL addresses, if this packet is to update a flow that is associated with addresses A and B, then it is also captured. That is, as long as addresses A and B that are embedded in the packet are matched, it is also captured.

For more information about clustering, see [Chapter 8, “Configuring a Cluster of ASAs.”](#)

Viewing the Crash Dump

If the ASA crashes, you can view the crash dump information. We recommend that you contact Cisco TAC if you want to interpret the crash dump. See the **show crashdump** command in the command reference.

Viewing the Coredump

A coredump is a snapshot of the running program when the program has terminated abnormally, or crashed. Coredumps are used to diagnose or debug errors and save a crash for future off-site analysis. Cisco TAC may request that you enable the coredump feature to troubleshoot application or system crashes on the ASA. See the **coredump** command in the command reference.



PART 9

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 44-1](#)
- [Licensing Requirements for Logging, page 44-5](#)
- [Prerequisites for Logging, page 44-5](#)
- [Guidelines and Limitations, page 44-6](#)
- [Configuring Logging, page 44-7](#)
- [Monitoring the Logs, page 44-20](#)
- [Configuration Examples for Logging, page 44-21](#)
- [Feature History for Logging, page 44-21](#)

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 44-2](#)
- [Analyzing Syslog Messages, page 44-2](#)
- [Syslog Message Format, page 44-3](#)
- [Severity Levels, page 44-3](#)
- [Message Classes and Range of Syslog IDs, page 44-4](#)
- [Filtering Syslog Messages, page 44-4](#)
- [Using Custom Message Lists, page 44-5](#)
- [Using Clustering, page 44-5](#)

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:

<i>ASA</i>	The syslog message facility code for messages that are generated by the ASA and ASASM. This value is always ASA.
<i>Level</i>	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 44-1 for more information.
<i>Message_number</i>	A unique six-digit number that identifies the syslog message.
<i>Message_text</i>	A text string that describes the condition. This portion of the syslog message sometimes includes IP addresses, port numbers, or usernames.

Severity Levels

[Table 44-1](#) 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 44-1 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.

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.

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, enter a new command 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 syslog messages 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 44-7](#)
- [Configuring an Output Destination, page 44-7](#)



Note

The minimum configuration depends on what you want to do and what your requirements are for handling syslog messages in the ASA and ASASM.

Enabling Logging

To enable logging, enter the following command:

Command	Purpose
<code>logging enable</code>	Enables logging. To disable logging, enter the no logging enable command.
Example: <code>hostname(config)# logging enable</code>	

What to Do Next

See the “[Configuring an Output Destination](#)” section on [page 44-7](#).

Configuring an Output Destination

To optimize syslog message usage for troubleshooting and performance monitoring, we recommend that you specify one or more locations where syslog messages should be sent, including an internal log buffer, one or more external syslog servers, ASDM, an SNMP management station, the console port, specified e-mail addresses, or Telnet and SSH sessions.

This section includes the following topics:

- [Sending Syslog Messages to an External Syslog Server, page 44-8](#)
- [Sending Syslog Messages to the Internal Log Buffer, page 44-9](#)
- [Sending Syslog Messages to an E-mail Address, page 44-11](#)
- [Sending Syslog Messages to ASDM, page 44-12](#)
- [Sending Syslog Messages to the Console Port, page 44-12](#)
- [Sending Syslog Messages to an SNMP Server, page 44-12](#)
- [Sending Syslog Messages to a Telnet or SSH Session, page 44-13](#)
- [Creating a Custom Event List, page 44-14](#)
- [Generating Syslog Messages in EMBLEM Format to a Syslog Server, page 44-15](#)
- [Generating Syslog Messages in EMBLEM Format to Other Output Destinations, page 44-15](#)
- [Changing the Amount of Internal Flash Memory Available for Logs, page 44-16](#)

- [Configuring the Logging Queue, page 44-16](#)
- [Sending All Syslog Messages in a Class to a Specified Output Destination, page 44-17](#)
- [Enabling Secure Logging, page 44-17](#)
- [Including the Device ID in Non-EMBLEM Format Syslog Messages, page 44-18](#)
- [Including the Date and Time in Syslog Messages, page 44-19](#)
- [Disabling a Syslog Message, page 44-19](#)
- [Changing the Severity Level of a Syslog Message, page 44-19](#)
- [Limiting the Rate of Syslog Message Generation, page 44-20](#)

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:

	Command	Purpose
Step 1	<pre>logging host interface_name syslog_ip [tcp[/port] udp[/port] [format emblem]]</pre> <p>Example: ciscoasa(config)# logging host dmz1 192.168.1.5 udp 1026 format emblem</p>	<p>Configures the ASA and ASASM to send messages to a syslog server.</p> <p>The format emblem keyword enables EMBLEM format logging for the syslog server with UDP only. The <i>interface_name</i> argument specifies the interface through which you access the syslog server. The <i>syslog_ip</i> argument specifies the IP address of the syslog server. The tcp[/port] or udp[/port] keyword and argument pair specify that the ASA and ASASM should use TCP or UDP to send syslog messages to the syslog server.</p> <p>You can configure the ASA 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.</p> <p>If you specify TCP, the ASA and ASASM discover when the syslog server fails and as a security protection, new connections through the ASA and ASA Services Module are blocked. To allow new connections regardless of connectivity to a TCP syslog server, see Step 3. If you specify UDP, the ASA and ASASM continue 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.</p>
Step 2	<pre>logging trap {severity_level message_list}</pre> <p>Example: ciscoasa(config)# logging trap errors</p>	<p>Specifies which syslog messages should be sent to the syslog server. You can specify the severity level number (1 through 7) or name. For example, if you set the severity level to 3, then the ASA and ASASM send syslog messages for severity levels 3, 2, and 1. You can specify a custom message list that identifies the syslog messages to send to the syslog server.</p>

	Command	Purpose
Step 3	logging permit-hostdown Example: ciscoasa(config)# logging permit-hostdown	(Optional) Disables the feature to block new connections when a TCP-connected syslog server is down. If the ASA or ASASM is configured to send syslog messages to a TCP-based syslog server, and if either the syslog server is down or the log queue is full, then new connections are blocked. New connections are allowed again after the syslog server is back up and the log queue is no longer full. For more information about the log queue, see the “Configuring the Logging Queue” section on page 44-16.
Step 4	logging facility <i>number</i> Example: ciscoasa(config)# logging facility 21	(Optional) Sets the logging facility to a value other than 20, which is what most UNIX systems expect.

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:

	Command	Purpose
Step 1	logging buffered { <i>severity_level</i> <i>message_list</i> } Example: ciscoasa(config)# logging buffered critical ciscoasa(config)# logging buffered level 2 ciscoasa(config)# logging buffered notif-list	Specifies which syslog messages should be sent to the internal log buffer, which serves as a temporary storage location. New messages are appended to the end of the list. When the buffer is full, that is, when the buffer wraps, old messages are overwritten as new messages are generated, unless you configure the ASA and ASASM to save the full buffer to another location. To empty the internal log buffer, enter the clear logging buffer command.
Step 2	logging buffer-size <i>bytes</i> Example: ciscoasa(config)# logging buffer-size 16384	Changes the size of the internal log buffer. The default buffer size is 4 KB.
Step 3	Choose one of the following options:	

Command	Purpose
<p>logging flash-bufferwrap</p> <p>Example: ciscoasa(config)# logging flash-bufferwrap</p>	<p>When saving the buffer content to another location, the ASA and ASASM create log files with names that use the following time-stamp format:</p> <p><i>LOG-YYYY-MM-DD-HHMMSS.TXT</i></p> <p>where <i>YYYY</i> is the year, <i>MM</i> is the month, <i>DD</i> is the day of the month, and <i>HHMMSS</i> is the time in hours, minutes, and seconds.</p> <p>The ASA and ASASM continue to save new messages to the internal log buffer and save the full log buffer content to the internal flash memory.</p>
<p>logging ftp-bufferwrap</p> <p>Example: ciscoasa(config)# logging ftp-bufferwrap</p>	<p>When saving the buffer content to another location, the ASA and ASASM create log files with names that use the following time-stamp format:</p> <p><i>LOG-YYYY-MM-DD-HHMMSS.TXT</i></p> <p>where <i>YYYY</i> is the year, <i>MM</i> is the month, <i>DD</i> is the day of the month, and <i>HHMMSS</i> is the time in hours, minutes, and seconds.</p> <p>The ASA and ASASM continue saving new messages to the internal log buffer and saves the full log buffer content to an FTP server.</p>
<p>logging ftp-server <i>server path username password</i></p> <p>Example: ciscoasa(config)# logging ftp-server 10.1.1.1 /syslogs logsupervisor 1luvMy10gs</p>	<p>Identifies the FTP server on which you want to store log buffer content. The <i>server</i> argument specifies the IP address of the external FTP server. The <i>path</i> argument specifies the directory path on the FTP server where the log buffer data is to be saved. This path is relative to the FTP root directory. The <i>username</i> argument specifies a username that is valid for logging into the FTP server. The <i>password</i> argument indicates the password for the username specified.</p>
<p>logging savefile [<i>savefile</i>]</p> <p>Example: ciscoasa(config)# logging savefile latest-logfile.txt</p>	<p>Saves the current log buffer content to the internal flash memory.</p>

Sending Syslog Messages to an E-mail Address

To send syslog messages to an e-mail address, perform the following steps:

	Command	Purpose
Step 1	logging mail { <i>severity_level</i> <i>message_list</i> } Example: hostname(config)# logging mail high-priority	Specifies which syslog messages should be sent to an e-mail address. When sent by e-mail, a syslog message appears in the subject line of the e-mail message. For this reason, we recommend configuring this option to notify administrators of syslog messages with high severity levels, such as critical, alert, and emergency.
Step 2	logging from-address <i>email_address</i> Example: ciscoasa(config)# logging from-address xxx-001@example.com	Specifies the source e-mail address to be used when sending syslog messages to an e-mail address.
Step 3	logging recipient-address <i>e-mail_address</i> [<i>severity_level</i>] Example: ciscoasa(config)# logging recipient-address admin@example.com	Specifies the recipient e-mail address to be used when sending syslog messages to an e-mail address.
Step 4	smtp-server <i>ip_address</i> Example: ciscoasa(config)# smtp-server 10.1.1.1	Specifies the SMTP server to be used when sending syslog messages to an e-mail address.

Sending Syslog Messages to ASDM

To send syslog messages to ASDM, perform the following steps:

	Command	Purpose
Step 1	<pre>logging asdm {severity_level message_list}</pre> <p>Example: ciscoasa(config)# logging asdm 2</p>	Specifies which syslog messages should be sent to ASDM. 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. Deletion of the oldest syslog message to make room for new ones is the default setting in ASDM. To control the number of syslog messages retained in the ASDM log buffer, you can change the size of the buffer.
Step 2	<pre>logging asdm-buffer-size num_of_msgs</pre> <p>Example: ciscoasa(config)# logging asdm-buffer-size 200</p>	Specifies the number of syslog messages to be retained in the ASDM log buffer. To empty the current content of the ASDM log buffer, enter the clear logging asdm command.

Sending Syslog Messages to the Console Port

To send syslog messages to the console port, enter the following command:

Command	Purpose
<pre>logging console {severity_level message_list}</pre> <p>Example: hostname(config)# logging console errors</p>	Specifies which syslog messages should be sent to the console port.

Sending Syslog Messages to an SNMP Server

To enable logging to an SNMP server, enter the following command:

Command	Purpose
<pre>logging history [logging_list level]</pre> <p>Example: hostname(config)# logging history errors</p>	Enables SNMP logging and specifies which messages are to be sent to SNMP servers. To disable SNMP logging, enter the no logging history command.

Sending Syslog Messages to a Telnet or SSH Session

To send syslog messages to a Telnet or SSH session, perform the following steps:

	Command	Purpose
Step 1	logging monitor { <i>severity_level</i> <i>message_list</i> } Example: ciscoasa(config)# logging monitor 6	Specifies which syslog messages should be sent to a Telnet or SSH session.
Step 2	terminal monitor Example: ciscoasa(config)# terminal monitor	Enables logging to the current session only. If you log out and then log in again, you need to reenter this command. To disable logging to the current session, enter the terminal no monitor command.

Creating a Custom Event List

To create a custom event list, perform the following steps:

Command	Purpose
<p>Step 1</p> <pre>logging list name {level level [class message_class] message start_id[-end_id]}</pre> <p>Example:</p> <pre>ciscoasa(config)# logging list notif-list level 3</pre>	<p>Specifies criteria for selecting messages to be saved in the internal log buffer. For example, if you set the severity level to 3, then the ASA sends syslog messages for severity levels 3, 2, and 1.</p> <p>The <i>name</i> argument specifies the name of the list. The level level keyword and argument pair specify the severity level. The class message_class keyword and argument pair specify a particular message class. The message start_id[-end_id] keyword and argument pair specify an individual syslog message number or a range of numbers.</p> <p>Note Do not use the names of severity levels as the name of a syslog message list. Prohibited names include emergencies, alert, critical, error, warning, notification, informational, and debugging. Similarly, do not use the first three characters of these words at the beginning of an event list name. For example, do not use an event list name that starts with the characters <i>err</i>.</p>
<p>Step 2</p> <pre>logging list name {level level [class message_class] message start_id[-end_id]}</pre> <p>Example:</p> <pre>ciscoasa(config)# logging list notif-list message 104024-105999</pre> <pre>ciscoasa(config)# logging list notif-list level critical</pre> <pre>ciscoasa(config)# logging list notif-list level warning class ha</pre>	<p>(Optional) Adds more criteria for message selection to the list. Enter the same command as in the previous step, specifying the name of the existing message list and the additional criterion. Enter a new command for each criterion that you want to add to the list. For example, you can specify criteria for syslog messages to be included in the list as the following:</p> <ul style="list-style-type: none"> • Syslog message IDs that fall into the range of 104024 to 105999. • All syslog messages with the critical severity level or higher (emergency, alert, or critical). • All ha class syslog messages with the warning severity level or higher (emergency, alert, critical, error, or warning). <p>Note A syslog message is logged if it satisfies any of these conditions. If a syslog message satisfies more than one of the conditions, the message is logged only once.</p>

Generating Syslog Messages in EMBLEM Format to a Syslog Server

To generate syslog messages in EMBLEM format to a syslog server, enter the following command:

Command	Purpose
<pre>logging host interface_name ip_address {tcp[/port] udp[/port]] [format emblem]</pre> <p>Example:</p> <pre>ciscoasa(config)# logging host interface_1 127.0.0.1 udp format emblem</pre>	<p>Sends syslog messages in EMBLEM format to a syslog server over UDP using port 514.</p> <p>The format emblem keyword enables EMBLEM format logging for the syslog server (UDP only). The <i>interface_name</i> argument specifies the interface through which you access the syslog server. The <i>ip_address</i> argument specifies the IP address of the syslog server. The tcp[/port] or udp[/port] keyword and argument pair specify that the ASA and ASASM should use TCP or UDP to send syslog messages to the syslog server.</p> <p>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.</p> <p>You can use multiple logging host commands to specify additional servers that would all receive syslog messages. If you configure two or more logging servers, make sure that you limit the logging severity level to warnings for all logging servers.</p> <p>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.</p> <p>Note Sending syslogs over TCP is not supported on a standby ASA.</p>

Generating Syslog Messages in EMBLEM Format to Other Output Destinations

To generate syslog messages in EMBLEM format to other output destinations, enter the following command:

Command	Purpose
<pre>logging emblem</pre> <p>Example:</p> <pre>ciscoasa(config)# logging emblem</pre>	<p>Sends syslog messages in EMBLEM format to output destinations other than a syslog server, such as Telnet or SSH sessions.</p>

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:

	Command	Purpose
Step 1	<pre>logging flash-maximum-allocation <i>kbytes</i></pre> <p>Example: <pre>hostname(config)# logging flash-maximum-allocation 1200</pre></p>	<p>Specifies the maximum amount of internal flash memory available for saving log files. By default, the ASA can use up to 1 MB of internal flash memory for log data. The minimum amount of internal flash memory that must be free for the ASA and ASASM to save log data is 3 MB.</p> <p>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.</p>
Step 2	<pre>logging flash-minimum-free <i>kbytes</i></pre> <p>Example: <pre>hostname(config)# logging flash-minimum-free 4000</pre></p>	<p>Specifies the minimum amount of internal flash memory that must be free for the ASA or ASASM to save a log file.</p>

Configuring the Logging Queue

To configure the logging queue, enter the following command:

Command	Purpose
<pre>logging queue <i>message_count</i></pre> <p>Example: <pre>ciscoasa(config)# logging queue 300</pre></p>	<p>Specifies the number of syslog messages that the ASA and ASASM can hold in its queue before sending them to the configured output destination. The ASA and ASASM have a fixed number of blocks in memory that can be allocated for buffering syslog messages while they are waiting to be sent to the configured output destination. The number of blocks required depends on the length of the syslog message queue and the number of syslog servers specified. The default queue size is 512 syslog messages. The queue size is limited only by block memory availability. Valid values are from 0 to 8192 messages, depending on the platform. If the logging queue is set to zero, the queue is the maximum configurable size (8192 messages), depending on the platform. The maximum queue size by platform is as follows:</p> <ul style="list-style-type: none"> • ASA-5505—1024 • ASA-5510—2048 • On all other platforms—8192

Sending All Syslog Messages in a Class to a Specified Output Destination

To send all syslog messages in a class to a specified output destination, enter the following command:

Command	Purpose
<pre>logging class message_class {buffered console history mail monitor trap} [severity_level]</pre> <p>Example: ciscoasa(config)# logging class ha buffered alerts</p>	<p>Overrides the configuration in the specified output destination command. For example, if you specify that messages at severity level 7 should go to the internal log buffer and that ha class messages at severity level 3 should go to the internal log buffer, then the latter configuration takes precedence. The buffered, history, mail, monitor, and trap keywords specify the output destination to which syslog messages in this class should be sent. The history keyword enables SNMP logging. The monitor keyword enables Telnet and SSH logging. The trap keyword enables syslog server logging. Select one destination per command line entry. To specify that a class should go to more than one destination, enter a new command for each output destination.</p>

Enabling Secure Logging

To enable secure logging, enter the following command:

Command	Purpose
<pre>logging host interface_name syslog_ip [tcp/port udp/port] [format emblem] [secure]</pre> <p>Example: ciscoasa(config)# logging host inside 10.0.0.1 TCP/1500 secure</p>	<p>Enables secure logging.</p> <p>The <i>interface_name</i> argument specifies the interface on which the syslog server resides. The <i>syslog_ip</i> argument specifies the IP address of the syslog server. The <i>port</i> argument specifies the port (TCP or UDP) that the syslog server listens to for syslog messages. The tcp keyword specifies that the ASA or ASASM should use TCP to send syslog messages to the syslog server. The udp keyword specifies that the ASA or ASASM should use UDP to send syslog messages to the syslog server. The format emblem keyword enables EMBLEM format logging for the syslog server. The secure keyword specifies that the connection to the remote logging host should use SSL/TLS for TCP only.</p> <p>Note Secure logging does not support UDP; an error occurs if you try to use this protocol.</p>

Including the Device ID in Non-EMBLEM Format Syslog Messages

To include the device ID in non-EMBLEM format syslog messages, enter the following command:

Command	Purpose
<p>logging device-id {cluster-id context-name hostname ipaddress <i>interface_name</i> [system] string <i>text</i>}</p> <p>Example:</p> <pre>ciscoasa(config)# logging device-id hostname</pre> <pre>ciscoasa(config)# logging device-id context-name</pre>	<p>Configures the ASA or ASASM to include a device ID in non-EMBLEM-format syslog messages. You can specify only one type of device ID for syslog messages. The context-name keyword indicates that the name of the current context should be used as the device ID (applies to multiple context mode only). If you enable the logging device ID for the admin context in multiple context mode, messages that originate in the system execution space use a device ID of system, and messages that originate in the admin context use the name of the admin context as the device ID.</p> <p>Note In an ASA cluster, always use the master unit IP address for the selected interface.</p> <p>The cluster-id keyword specifies the unique name in the boot configuration of an individual ASA unit in the cluster as the device ID. The hostname keyword specifies that the hostname of the ASA should be used as the device ID. The ipaddress <i>interface_name</i> keyword-argument pair specifies that the interface IP address specified as <i>interface_name</i> should be used as the device ID. If you use the ipaddress keyword, the device ID becomes the specified ASA interface IP address, regardless of the interface from which the syslog message is sent. In the cluster environment, the system keyword dictates that the device ID becomes the system IP address on the interface. This keyword provides a single, consistent device ID for all syslog messages that are sent from the device. The string <i>text</i> keyword-argument pair specifies that the text string should be used as the device ID. The string can include as many as 16 characters.</p> <p>You cannot use blank spaces or any of the following characters:</p> <ul style="list-style-type: none"> • & (ampersand) • ‘ (single quote) • “ (double quote) • < (less than) • > (greater than) • ? (question mark) <p>Note If enabled, the device ID does not appear in EMBLEM-formatted syslog messages nor in SNMP traps.</p>

Including the Date and Time in Syslog Messages

To include the date and time in syslog messages, enter the following command:

Command	Purpose
logging timestamp ciscoasa(config)# logging timestamp Example: ciscoasa(config)# logging timestamp LOG-2008-10-24-081856.TXT	Specifies that syslog messages should include the date and time that they were generated. To remove the date and time from syslog messages, enter the no logging timestamp command.

Disabling a Syslog Message

To disable a specified syslog message, enter the following command:

Command	Purpose
no logging message <i>message_number</i> Example: ciscoasa(config)# no logging message 113019	Prevents the ASA or ASASM from generating a particular syslog message. To reenab a disabled syslog message, enter the logging message <i>message_number</i> command (for example, logging message 113019). To reenab logging of all disabled syslog messages, enter the clear config logging disabled command.

Changing the Severity Level of a Syslog Message

To change the severity level of a syslog message, enter the following command:

Command	Purpose
logging message <i>message_ID</i> level <i>severity_level</i> Example: hostname(config)# logging message 113019 level 5	Specifies the severity level of a syslog message. To reset the severity level of a syslog message to its setting, enter the no logging message <i>message_ID</i> level <i>current_severity_level</i> command (for example, no logging message 113019 level 5). To reset the severity level of all modified syslog messages to their settings, enter the clear configure logging level command.

Limiting the Rate of Syslog Message Generation

To limit the rate of syslog message generation, enter the following command:

Command	Purpose
<pre>logging rate-limit {unlimited {num [interval]}}</pre> <pre>message <i>syslog_id</i> level <i>severity_level</i></pre> <p>Example:</p> <pre>hostname(config)# logging rate-limit 1000 600 level 6</pre>	<p>Applies a specified severity level (1 through 7) to a set of messages or to an individual message (not the destination) within a specified time period. Rate limits affect the volume of messages being sent to all configured destinations. To reset the logging rate limit to the default value, enter the clear running-config logging rate-limit command. To reset the logging rate limit, enter the clear configure logging rate-limit command.</p>

Monitoring the Logs

To monitor the logs in the log buffer or in real-time and assist in monitoring the system performance, enter one of the following commands:

Command	Purpose
show logging	Shows syslog messages, including the severity level. Note The maximum number of syslog messages that are available to view is 1000, which is the default setting. The maximum number of syslog messages that are available to view is 2000.
show logging message	Shows a list of syslog messages with modified severity levels and disabled syslog messages.
show logging message <i>message_ID</i>	Shows the severity level of a specific syslog message.
show logging queue	Shows the logging queue and queue statistics.
show logging rate-limit	Shows the disallowed syslog messages.
show running-config logging rate-limit	Shows the current logging rate-limit setting.

Examples

The following example shows the logging information that displays for the **show logging** command:

```
hostname(config)# show logging
Syslog logging: enabled
  Facility: 16
  Timestamp logging: disabled
  Standby logging: disabled
  Deny Conn when Queue Full: disabled
  Console logging: disabled
  Monitor logging: disabled
  Buffer logging: disabled
  Trap logging: level errors, facility 16, 3607 messages logged
    Logging to infrastructure 10.1.2.3
  History logging: disabled
  Device ID: 'inside' interface IP address "10.1.1.1"
```

```
Mail logging: disabled
ASDM logging: disabled
```

Configuration Examples for Logging

The following examples show how to control both whether a syslog message is enabled and the severity level of the specified syslog message:

```
ciscoasa(config)# show logging message 403503
syslog 403503: -level errors (enabled)

ciscoasa(config)# logging message 403503 level 1
ciscoasa(config)# show logging message 403503
syslog 403503: -level errors, current-level alerts (enabled)

ciscoasa(config)# no logging message 403503
ciscoasa(config)# show logging message 403503
syslog 403503: -level errors, current-level alerts (disabled)

ciscoasa(config)# logging message 403503
ciscoasa(config)# show logging message 403503
syslog 403503: -level errors, current-level alerts (enabled)

ciscoasa(config)# no logging message 403503 level 3
ciscoasa(config)# show logging message 403503
syslog 403503: -level errors (enabled)
```

Feature History for Logging

Table 44-2 lists each feature change and the platform release in which it was implemented.

Table 44-2 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.
Rate limit	7.0(4)	Limits the rate at which syslog messages are generated. We introduced the following command: 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 introduced the following command: logging list .

Table 44-2 Feature History for Logging (continued)

Feature Name	Platform Releases	Feature Information
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 command: logging host .
Logging class	8.0(4), 8.1(1)	Added support for the ipaa event class of logging messages. We modified the following command: logging class .
Logging class and saved logging buffers	8.2(1)	Added support for the dap event class of logging messages. We modified the following command: logging class . Added support to clear the saved logging buffers (ASDM, internal, FTP, and flash). We introduced the following command: clear logging queue bufferwrap .
Password encryption	8.3(1)	Added support for password encryption. We modified the following command: logging ftp server .
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 use the logging permit-hostdown command. We modified the following command: show logging . We introduced the following syslog messages: 414005, 414006, 414007, and 414008.
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 command: logging device-id .



Configuring SNMP

This chapter describes how to configure SNMP to monitor the ASA and ASASM and includes the following sections:

- [Information About SNMP, page 45-1](#)
- [Licensing Requirements for SNMP, page 45-17](#)
- [Prerequisites for SNMP, page 45-17](#)
- [Guidelines and Limitations, page 45-17](#)
- [Configuring SNMP, page 45-18](#)
- [Troubleshooting Tips, page 45-24](#)
- [Monitoring SNMP, page 45-26](#)
- [Configuration Examples for SNMP, page 45-28](#)
- [Where to Go Next, page 45-29](#)
- [Additional References, page 45-29](#)
- [Feature History for SNMP, page 45-31](#)

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 45-2](#)
- [Information About MIBs and Traps, page 45-3](#)
- [SNMP Object Identifiers, page 45-3](#)
- [SNMP Physical Vendor Type Values, page 45-5](#)
- [Supported Tables in MIBs, page 45-11](#)
- [Supported Traps \(Notifications\), page 45-12](#)
- [SNMP Version 3, page 45-15](#)

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 45-1 lists the terms that are commonly used when working with SNMP:

Table 45-1 SNMP Terminology

Term	Description
Agent	The SNMP server running on the ASA. The SNMP agent has the following features: <ul style="list-style-type: none"> • 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.

Information About MIBs and Traps

MIBs are either standard or enterprise-specific. Standard MIBs are created by the IETF and documented in various RFCs. A trap reports significant events occurring on a network device, most often errors or failures. SNMP traps are defined in either standard or enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. SNMP traps are compiled into the ASA or ASASM software.

If needed, you can also download RFCs, standard MIBs, and standard traps from the following locations:

<http://www.ietf.org/>

<ftp://ftp-sj.cisco.com/pub/mibs>

Download a complete list of Cisco MIBs, traps, and OIDs from the following location:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

In addition, download Cisco OIDs by FTP from the following location:

<ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz>



Note

In software versions 7.2(1), 8.0(2), and later, the interface information accessed through SNMP refreshes about every 5 seconds. As a result, we recommend that you wait for at least 5 seconds between consecutive polls.

SNMP Object Identifiers

Each Cisco system-level product has an SNMP object identifier (OID) for use as a MIB-II sysObjectID. The CISCO-PRODUCTS-MIB includes the OIDs that can be reported in the sysObjectID object in the SNMPv2-MIB. You can use this value to identify the model type. [Table 45-2](#) lists the sysObjectID OIDs for ASA models.

Table 45-2 *SNMP Object Identifiers*

Product Identifier	sysObjectID	Model Number
ASA 5505	ciscoASA5505 (ciscoProducts 745)	Cisco ASA 5505
ASA 5510	ciscoASA5510 (ciscoProducts 669)	Cisco ASA 5510
ASA 5510	ciscoASA5510sc (ciscoProducts 773)	Cisco ASA 5510 security context
ASA 5510	ciscoASA5510sy (ciscoProducts 774)	Cisco ASA 5510 system context
ASA 5520	ciscoASA5520 (ciscoProducts 670)	Cisco ASA 5520
ASA 5520	ciscoASA5520sc (ciscoProducts 671)	Cisco ASA 5520 security context
ASA 5520	ciscoASA5520sy (ciscoProducts 764)	Cisco ASA 5520 system context
ASA 5540	ciscoASA5540 (ciscoProducts 672)	Cisco ASA 5540
ASA 5540	ciscoASA5540sc (ciscoProducts 673)	Cisco ASA 5540 security context
ASA 5540	ciscoASA5540sy (ciscoProducts 765)	Cisco ASA 5540 system context
ASA 5550	ciscoASA5550 (ciscoProducts 753)	Cisco ASA 5550
ASA 5550	ciscoASA5550sc (ciscoProducts 763)	Cisco ASA 5550 security context
ASA 5550	ciscoASA 5550sy (ciscoProducts 766)	Cisco ASA 5550 system context

Table 45-2 SNMP Object Identifiers (continued)

ASA5580	ciscoASA5580 (ciscoProducts 914)	Cisco ASA 5580
ASA5580	ciscoASA5580 (ciscoProducts 915)	Cisco ASA 5580 security context
ASA5580	ciscoASA5580 (ciscoProducts 916)	Cisco ASA 5580 system context
ASA5585-SSP10	ciscoASA5585Ssp10 (ciscoProducts 1194)	ASA 5585-X SSP-10
ASA5585-SSP20	ciscoASA5585Ssp20 (ciscoProducts 1195)	ASA 5585-X SSP-20
ASA5585-SSP40	ciscoASA5585Ssp40 (ciscoProducts 1196)	ASA 5585-X SSP-40
ASA5585-SSP60	ciscoASA5585Ssp60 (ciscoProducts 1197)	ASA 5585-X SSP-60
ASA5585-SSP10	ciscoASA5585Ssp10sc (ciscoProducts 1198)	ASA 5585-X SSP-10 security context
ASA5585-SSP20	ciscoASA5585Ssp20sc (ciscoProducts 1199)	ASA 5585-X SSP-20 security context
ASA5585-SSP40	ciscoASA5585Ssp40sc (ciscoProducts 1200)	ASA 5585-X SSP-40 security context
ASA5585-SSP60	ciscoASA5585Ssp60sc (ciscoProducts 1201)	ASA 5585-X SSP-60 security context
ASA5585-SSP10	ciscoASA5585Ssp10sy (ciscoProducts 1202)	ASA 5585-X SSP-10 system context
ASA5585-SSP20	ciscoASA5585Ssp20sy (ciscoProducts 1203)	ASA 5585-X SSP-20 system context
ASA5585-SSP40	ciscoASA5585Ssp40sy (ciscoProducts 1204)	ASA 5585-X SSP-40 system context
ASA5585-SSP60	ciscoASA5585Ssp60sy (ciscoProducts 1205)	ASA 5585-X SSP-60 system context
ASA Services Module for Catalyst switches/7600 routers	ciscoAaSm1 (ciscoProducts 1277)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers
ASA Services Module for Catalyst switches/7600 routers security context	ciscoAaSm1sc (ciscoProducts 1275)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers security context
ASA Services Module for Catalyst switches/7600 routers security context with No Payload Encryption	ciscoAaSm1K7sc (ciscoProducts 1334)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers security context with No Payload Encryption
ASA Services Module for Catalyst switches/7600 routers system context	ciscoAaSm1sy (ciscoProducts 1276)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers system context
ASA Services Module for Catalyst switches system context/7600 routers with No Payload Encryption	ciscoAaSm1K7sy (ciscoProducts 1335)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers system context with No Payload Encryption
ASA Services Module for Catalyst switches/7600 routers system context with No Payload Encryption	ciscoAaSm1K7 (ciscoProducts 1336)	Adaptive Security Appliance (ASA) Services Module for Catalyst switches/7600 routers with No Payload Encryption
ASA 5512	ciscoASA5512 (ciscoProducts 1407)	ASA 5512 Adaptive Security Appliance
ASA 5525	ciscoASA5525 (ciscoProducts 1408)	ASA 5525 Adaptive Security Appliance
ASA 5545	ciscoASA5545 (ciscoProducts 1409)	ASA 5545 Adaptive Security Appliance
ASA 5555	ciscoASA5555 (ciscoProducts 1410)	ASA 5555 Adaptive Security Appliance

Table 45-2 *SNMP Object Identifiers (continued)*

ASA 5512 Security Context	ciscoASA5512sc (ciscoProducts 1411)	ASA 5512 Adaptive Security Appliance Security Context
ASA 5525 Security Context	ciscoASA5525sc (ciscoProducts 1412)	ASA 5525 Adaptive Security Appliance Security Context
ASA 5545 Security Context	ciscoASA5545sc (ciscoProducts 1413)	ASA 5545 Adaptive Security Appliance Security Context
ASA 5555 Security Context	ciscoASA5555sc (ciscoProducts 1414)	ASA 5555 Adaptive Security Appliance Security Context
ASA 5512 System Context	ciscoASA5512sy (ciscoProducts 1415)	ASA 5512 Adaptive Security Appliance System Context
ASA 5515 System Context	ciscoASA5515sy (ciscoProducts 1416)	ASA 5515 Adaptive Security Appliance System Context
ASA 5525 System Context	ciscoASA5525sy (ciscoProducts 1417)	ASA 5525 Adaptive Security Appliance System Context
ASA 5545 System Context	ciscoASA5545sy (ciscoProducts 1418)	ASA 5545 Adaptive Security Appliance System Context
ASA 5555 System Context	ciscoASA5555sy (ciscoProducts 1419)	ASA 5555 Adaptive Security Appliance System Context
ASA 5515 Security Context	ciscoASA5515sc (ciscoProducts 1420)	ASA 5515 Adaptive Security Appliance System Context
ASA 5515	ciscoASA5515 (ciscoProducts 1421)	ASA 5515 Adaptive Security Appliance

SNMP Physical Vendor Type Values

Each Cisco chassis or standalone system has a unique type number for SNMP use. The entPhysicalVendorType OIDs are defined in the CISCO-ENTITY-VENDORTYPE-OID-MIB. This value is returned in the entPhysicalVendorType object from the ASA or ASASM SNMP agent. You can use this value to identify the type of component (module, power supply, fan, sensors, CPU, and so on). [Table 45-3](#) lists the physical vendor type values for the ASA and ASASM models.

Table 45-3 *SNMP Physical Vendor Type Values*

Item	entPhysicalVendorType OID Description
ASA Services Module for Catalyst switches/7600 routers	cevCat6kWsSvcAsaSm1 (cevModuleCat6000Type 169)
ASA Services Module for Catalyst switches/7600 routers with No Payload Encryption	cevCat6kWsSvcAsaSm1K7 (cevModuleCat6000Type 186)
ASA 5505 chassis	cevChassisASA5505 (cevChassis 560)
ASA 5510 chassis	cevChassisASA5510 (cevChassis 447)
Cisco Adaptive Security Appliance (ASA) 5512 Adaptive Security Appliance	cevChassisASA5512 (cevChassis 1113)
Cisco Adaptive Security Appliance (ASA) 5512 Adaptive Security Appliance with No Payload Encryption	cevChassisASA5512K7 (cevChassis 1108)

Table 45-3 *SNMP Physical Vendor Type Values (continued)*

Cisco Adaptive Security Appliance (ASA) 5515 Adaptive Security Appliance	cevChassisASA5515 (cevChassis 1114)
Cisco Adaptive Security Appliance (ASA) 5515 Adaptive Security Appliance with No Payload Encryption	cevChassisASA5515K7 (cevChassis 1109)
ASA 5520 chassis	cevChassisASA5520 (cevChassis 448)
Cisco Adaptive Security Appliance (ASA) 5525 Adaptive Security Appliance	cevChassisASA5525 (cevChassis 1115)
Cisco Adaptive Security Appliance (ASA) 5525 Adaptive Security Appliance with No Payload Encryption	cevChassisASA5525K7 (cevChassis 1110)
ASA 5540 chassis	cevChassisASA5540 (cevChassis 449)
Cisco Adaptive Security Appliance (ASA) 5545 Adaptive Security Appliance	cevChassisASA5545 (cevChassis 1116)
Cisco Adaptive Security Appliance (ASA) 5545 Adaptive Security Appliance with No Payload Encryption	cevChassisASA5545K7 (cevChassis 1111)
ASA 5550 chassis	cevChassisASA5550 (cevChassis 564)
Cisco Adaptive Security Appliance (ASA) 5555 Adaptive Security Appliance	cevChassisASA5555 (cevChassis 1117)
Cisco Adaptive Security Appliance (ASA) 5555 Adaptive Security Appliance with No Payload Encryption	cevChassisASA5555K7 (cevChassis 1112)
ASA 5580 chassis	cevChassisASA5580 (cevChassis 704)
Central Processing Unit for Cisco Adaptive Security Appliance 5512	cevCpuAsa5512 (cevModuleCpuType 229)
Central Processing Unit for Cisco Adaptive Security Appliance 5512 with no Payload Encryption	cevCpuAsa5512K7 (cevModuleCpuType 224)
Central Processing Unit for Cisco Adaptive Security Appliance 5515	cevCpuAsa5515 (cevModuleCpuType 230)
Central Processing Unit for Cisco Adaptive Security Appliance 5515 with no Payload Encryption	cevCpuAsa5515K7 (cevModuleCpuType 225)
Central Processing Unit for Cisco Adaptive Security Appliance 5525	cevCpuAsa5525 (cevModuleCpuType 231)
Central Processing Unit for Cisco Adaptive Security Appliance 5525 with no Payload Encryption	cevCpuAsa5525K7 (cevModuleCpuType 226)
Central Processing Unit for Cisco Adaptive Security Appliance 5545	cevCpuAsa5545 (cevModuleCpuType 232)
Central Processing Unit for Cisco Adaptive Security Appliance 5545 with no Payload Encryption	cevCpuAsa5545K7 (cevModuleCpuType 227)
Central Processing Unit for Cisco Adaptive Security Appliance 5555	cevCpuAsa5555 (cevModuleCpuType 233)

Table 45-3 *SNMP Physical Vendor Type Values (continued)*

Central Processing Unit for Cisco Adaptive Security Appliance 5555 with no Payload Encryption	cevCpuAsa5555K7 (cevModuleCpuType 228)
CPU for ASA 5580	cevCpuAsa5580 (cevModuleType 200)
CPU for ASA 5585 SSP-10	cevCpuAsa5585Ssp10 (cevModuleCpuType 204)
CPU for ASA 5585 SSP-10 No Payload Encryption	cevCpuAsa5585Ssp10K7 (cevModuleCpuType 205)
CPU for ASA 5585 SSP-20	cevCpuAsa5585Ssp20 (cevModuleCpuType 206)
CPU for ASA 5585 SSP-20 No Payload Encryption	cevCpuAsa5585Ssp20K7 (cevModuleCpuType 207)
CPU for ASA 5585 SSP-40	cevCpuAsa5585Ssp40 (cevModuleCpuType 208)
CPU for ASA 5585 SSP-40 No Payload Encryption	cevCpuAsa5585Ssp40K7 (cevModuleCpuType 209)
CPU for ASA 5585 SSP-60	cevCpuAsa5585Ssp60 (cevModuleCpuType 210)
CPU for ASA 5585 SSP-60 No Payload Encryption	cevCpuAsa5585Ssp60K7 (cevModuleCpuType 211)
CPU for Cisco ASA Services Module for Catalyst switches/7600 routers	cevCpuAsaSm1 (cevModuleCpuType 222)
CPU for Cisco ASA Services Module with No Payload Encryption for Catalyst switches/7600 routers	cevCpuAsaSm1K7 (cevModuleCpuType 223)
Chassis Cooling Fan in Adaptive Security Appliance 5512	cevFanASA5512ChassisFan (cevFan 163)
Chassis Cooling Fan in Adaptive Security Appliance 5512 with No Payload Encryption	cevFanASA5512K7ChassisFan (cevFan 172)
Chassis Cooling Fan in Adaptive Security Appliance 5515	cevFanASA5515ChassisFan (cevFan 164)
Chassis Cooling Fan in Adaptive Security Appliance 5515 with No Payload Encryption	cevFanASA5515K7ChassisFan (cevFan 171)
Chassis Cooling Fan in Adaptive Security Appliance 5525	cevFanASA5525ChassisFan (cevFan 165)
Chassis Cooling Fan in Adaptive Security Appliance 5525 with No Payload Encryption	cevFanASA5525K7ChassisFan (cevFan 170)
Chassis Cooling Fan in Adaptive Security Appliance 5545	cevFanASA5545ChassisFan (cevFan 166)
Chassis Cooling Fan in Adaptive Security Appliance 5545 with No Payload Encryption	cevFanASA5545K7ChassisFan (cevFan 169)
Power Supply Fan in Adaptive Security Appliance 5545 with No Payload Encryption	cevFanASA5545K7PSFan (cevFan 161)
Power Supply Fan in Adaptive Security Appliance 5545	cevFanASA5545PSFan (cevFan 159)
Chassis Cooling Fan in Adaptive Security Appliance 5555	cevFanASA5555ChassisFan (cevFan 167)
Chassis Cooling Fan in Adaptive Security Appliance 5555 with No Payload Encryption	cevFanASA5555K7ChassisFan (cevFan 168)
Power Supply Fan in Adaptive Security Appliance 5555	cevFanASA5555PSFan (cevFan 160)

Table 45-3 *SNMP Physical Vendor Type Values (continued)*

Power Supply Fan in Adaptive Security Appliance 5555 with No Payload Encryption	cevFanASA5555PSFanK7 (cevFan 162)
Fan type for ASA 5580	cevFanASA5580Fan (cevFan 138)
Power supply fan for ASA 5585-X	cevFanASA5585PSFan (cevFan 146)
ASA 5580 4-port GE copper interface card	cevModuleASA5580Pm4xlgeCu (cevModuleASA5580Type 1)
10-Gigabit Ethernet interface	cevPort10GigEthernet (cevPort 315)
Gigabit Ethernet port	cevPortGe (cevPort 109)
Power Supply unit in Adaptive Security Appliance 5545	cevPowerSupplyASA5545PSInput (cevPowerSupply 323)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5545	cevPowerSupplyASA5545PSPresence (cevPowerSupply 321)
Power Supply unit in Adaptive Security Appliance 5555	cevPowerSupplyASA5555PSInput (cevPowerSupply 324)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5555	cevPowerSupplyASA5555PSPresence (cevPowerSupply 322)
Power supply input for ASA 5580	cevPowerSupplyASA5580PSInput (cevPowerSupply 292)
Power supply input for ASA 5585	cevPowerSupplyASA5585PSInput (cevPowerSupply 304)
Cisco Adaptive Security Appliance (ASA) 5512 Chassis Fan sensor	cevSensorASA5512ChassisFanSensor (cevSensor 120)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5512	cevSensorASA5512ChassisTemp (cevSensor 107)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5512	cevSensorASA5512CPUTemp (cevSensor 96)
Cisco Adaptive Security Appliance (ASA) 5512 with No Payload Encryption Chassis Fan sensor	cevSensorASA5512K7ChassisFanSensor (cevSensor 125)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5512 with No Payload Encryption	cevSensorASA5512K7CPUTemp (cevSensor 102)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5512 with No Payload Encryption	cevSensorASA5512K7PSFanSensor (cevSensor 116)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5512	cevSensorASA5512PSFanSensor (cevSensor 119)
Cisco Adaptive Security Appliance (ASA) 5515 Chassis Fan sensor	cevSensorASA5515ChassisFanSensor (cevSensor 121)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5515	cevSensorASA5515ChassisTemp (cevSensor 98)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5515	cevSensorASA5515CPUTemp (cevSensor 97)
Cisco Adaptive Security Appliance (ASA) 5515 with No Payload Encryption Chassis Fan sensor	cevSensorASA5515K7ChassisFanSensor (cevSensor 126)

Table 45-3 *SNMP Physical Vendor Type Values (continued)*

Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5515 with No Payload Encryption	cevSensorASA5515K7CPUTemp (cevSensor 103)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5515 with No Payload Encryption	cevSensorASA5515K7PSFanSensor (cevSensor 115)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5515	cevSensorASA5515PSFanSensor (cevSensor 118)
Cisco Adaptive Security Appliance (ASA) 5525 Chassis Fan sensor	cevSensorASA5525ChassisFanSensor (cevSensor 122)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5525	cevSensorASA5525ChassisTemp (cevSensor 108)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5525	cevSensorASA5525CPUTemp (cevSensor 99)
Cisco Adaptive Security Appliance (ASA) 5525 with No Payload Encryption Chassis Fan sensor	cevSensorASA5525K7ChassisFanSensor (cevSensor 127)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5525 with No Payload Encryption	cevSensorASA5525K7CPUTemp (cevSensor 104)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5525 with No Payload Encryption	cevSensorASA5525K7PSFanSensor (cevSensor 114)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5525	cevSensorASA5525PSFanSensor (cevSensor 117)
Cisco Adaptive Security Appliance (ASA) 5545 Chassis Fan sensor	cevSensorASA5545ChassisFanSensor (cevSensor 123)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5545	cevSensorASA5545ChassisTemp (cevSensor 109)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5545	cevSensorASA5545CPUTemp (cevSensor 100)
Cisco Adaptive Security Appliance (ASA) 5545 with No Payload Encryption Chassis Fan sensor	cevSensorASA5545K7ChassisFanSensor (cevSensor 128)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545K7ChassisTemp (cevSensor 90)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545K7CPUTemp (cevSensor 105)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545K7PSFanSensor (cevSensor 113)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545K7PSPresence (cevSensor 87)
Temperature Sensor for Power Supply Fan in Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545K7PSTempSensor (cevSensor 94)

Table 45-3 *SNMP Physical Vendor Type Values (continued)*

Sensor for Power Supply Fan in Adaptive Security Appliance 5545 with No Payload Encryption	cevSensorASA5545PSFanSensor (cevSensor 89)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5545	cevSensorASA5545PSPresence (cevSensor 130)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5555	cevSensorASA5545PSPresence (cevSensor 131)
Temperature Sensor for Power Supply Fan in Adaptive Security Appliance 5545	cevSensorASA5545PSTempSensor (cevSensor 92)
Cisco Adaptive Security Appliance (ASA) 5555 Chassis Fan sensor	cevSensorASA5555ChassisFanSensor (cevSensor 124)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5555	cevSensorASA5555ChassisTemp (cevSensor 110)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5555	cevSensorASA5555CPUTemp (cevSensor 101)
Cisco Adaptive Security Appliance (ASA) 5555 with No Payload Encryption Chassis Fan sensor	cevSensorASA5555K7ChassisFanSensor (cevSensor 129)
Chassis Ambient Temperature Sensor for Cisco Adaptive Security Appliance 5555 with No Payload Encryption	cevSensorASA5555K7ChassisTemp (cevSensor 111)
Central Processing Unit Temperature Sensor for Cisco Adaptive Security Appliance 5555 with No Payload Encryption	cevSensorASA5555K7CPUTemp (cevSensor 106)
Sensor for Chassis Cooling Fan in Adaptive Security Appliance 5555 with No Payload Encryption	cevSensorASA5555K7PSFanSensor (cevSensor 112)
Presence Sensor for Power Supply input in Adaptive Security Appliance 5555 with No Payload Encryption	cevSensorASA5555K7PSPresence (cevSensor 88)
Temperature Sensor for Power Supply Fan in Adaptive Security Appliance 5555 with No Payload Encryption	cevSensorASA5555K7PSTempSensor (cevSensor 95)
Sensor for Power Supply Fan in Adaptive Security Appliance 5555	cevSensorASA5555PSFanSensor (cevSensor 91)
Temperature Sensor for Power Supply Fan in Adaptive Security Appliance 5555	cevSensorASA5555PSTempSensor (cevSensor 93)
Sensor type for ASA 5580	cevSensorASA5580FanSensor (cevSensor 76)
Sensor for power supply input for ASA 5580	cevSensorASA5580PSInput (cevSensor 74)
Sensor for power supply fan for ASA 5585-X	cevSensorASA5585PSFanSensor (cevSensor 86)
Sensor for power supply input for ASA 5585-X	cevSensorASA5585PSInput (cevSensor 85)
CPU temperature sensor for ASA 5585 SSP-10	cevSensorASA5585SSp10CPUTemp (cevSensor 77)
CPU temperature sensor for ASA 5585 SSP-10 No Payload Encryption	cevSensorASA5585SSp10K7CPUTemp (cevSensor 78)
CPU temperature sensor for ASA 5585 SSP-20	cevSensorASA5585SSp20CPUTemp (cevSensor 79)
CPU temperature sensor for ASA 5585 SSP-20 No Payload Encryption	cevSensorASA5585SSp20K7CPUTemp (cevSensor 80)

Table 45-3 SNMP Physical Vendor Type Values (continued)

CPU temperature sensor for ASA 5585 SSP-40	cevSensorASA5585SSp40CPUTemp (cevSensor 81)
CPU temperature sensor for ASA 5585 SSP-40 No Payload Encryption	cevSensorASA5585SSp40K7CPUTemp (cevSensor 82)
CPU temperature sensor for ASA 5585 SSP-60	cevSensorASA5585SSp60CPUTemp (cevSensor 83)
CPU temperature sensor for ASA 5585 SSP-60 No Payload Encryption	cevSensorASA5585SSp60K7CPUTemp (cevSensor 84)
Adaptive Security Appliance 5555-X Field-Replaceable Solid State Drive	cevModuleASA5555XFRSSD (cevModuleCommonCards 396)
Adaptive Security Appliance 5545-X Field-Replaceable Solid State Drive	cevModuleASA5545XFRSSD (cevModuleCommonCards 397)
Adaptive Security Appliance 5525-X Field-Replaceable Solid State Drive	cevModuleASA5525XFRSSD (cevModuleCommonCards 398)
Adaptive Security Appliance 5515-X Field-Replaceable Solid State Drive	cevModuleASA5515XFRSSD (cevModuleCommonCards 399)
Adaptive Security Appliance 5512-X Field-Replaceable Solid State Drive	cevModuleASA5512XFRSSD (cevModuleCommonCards 400)

Supported Tables in MIBs

Table 45-4 lists the supported tables and objects for the specified MIBs.

Table 45-4 Supported Tables and Objects in MIBs

MIB Name	Supported Tables and Objects
CISCO-ENHANCED-MEMPOOL-MIB	compMemPoolTable, compMemPoolIndex, compMemPoolType, compMemPoolName, compMemPoolAlternate, compMemPoolValid, compMemPoolUsed, compMemPoolFree, compMemPoolUsedOvrflw, compMemPoolHCUsed, compMemPoolFreeOvrflw, compMemPoolHCFree
CISCO-ENTITY-SENSOR-EXT-MIB Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.	ceSensorExtThresholdTable
CISCO-L4L7MODULE-RESOURCE-LIMIT-MIB	ciscoL4L7ResourceLimitTable
CISCO-TRUSTSEC-SXP-MIB	ctsXspGlobalObjects, ctsXspConnectionObjects, ctsXspSgtObjects
DISMAN-EVENT-MIB	mteTriggerTable, mteTriggerThresholdTable, mteObjectsTable, mteEventTable, mteEventNotificationTable
DISMAN-EXPRESSION-MIB Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.	expExpressionTable, expObjectTable, expValueTable

Table 45-4 Supported Tables and Objects in MIBs (continued)

ENTITY-SENSOR-MIB Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.	entPhySensorTable
NAT-MIB	natAddrMapTable, natAddrMapIndex, natAddrMapName, natAddrMapGlobalAddrType, natAddrMapGlobalAddrFrom, natAddrMapGlobalAddrTo, natAddrMapGlobalPortFrom, natAddrMapGlobalPortTo, natAddrMapProtocol, natAddrMapAddrUsed, natAddrMapRowStatus

Supported Traps (Notifications)

Table 45-5 lists the supported traps (notifications) and their associated MIBs.

Table 45-5 Supported Traps (Notifications)

Trap and MIB Name	Varbind List	Description
authenticationFailure (SNMPv2-MIB)	—	For SNMP Version 1 or 2, the community string provided in the SNMP request is incorrect. For SNMP Version 3, a report PDU is generated instead of a trap if the auth or priv passwords or usernames are incorrect. The snmp-server enable traps snmp authentication command is used to enable and disable transmission of these traps.
cefcFRUInserted (CISCO-ENTITY-FRU-CONTROL-MIB)	—	The snmp-server enable traps entity fru-insert command is used to enable this notification.
cefcFRURemoved (CISCO-ENTITY-FRU-CONTROL-MIB)	—	The snmp-server enable traps entity fru-remove command is used to enable this notification.

Table 45-5 Supported Traps (Notifications) (continued)

<p>ceSensorExtThresholdNotification (CISCO-ENTITY-SENSOR-EXT-MIB)</p> <p>Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.</p>	<p>ceSensorExtThresholdValue, entPhySensorValue, entPhySensorType, entPhysicalName</p>	<p>The snmp-server enable traps entity [power-supply-failure fan-failure cpu-temperature] command is used to enable transmission of the entity threshold notifications. This notification is sent for a power supply failure. The objects sent identify the fan and CPU temperature.</p> <p>The snmp-server enable traps entity fan-failure command is used to enable transmission of the fan failure trap.</p> <p>The snmp-server enable traps entity power-supply-failure command is used to enable transmission of the power supply failure trap.</p> <p>The snmp-server enable traps entity chassis-fan-failure command is used to enable transmission of the chassis fan failure trap.</p> <p>The snmp-server enable traps entity cpu-temperature command is used to enable transmission of the high CPU temperature trap.</p> <p>The snmp-server enable traps entity power-supply-presence command is used to enable transmission of the power supply presence failure trap.</p> <p>The snmp-server enable traps entity power-supply-temperature command is used to enable transmission of the power supply temperature threshold trap.</p> <p>The snmp-server enable traps entity chassis-temperature command is used to enable transmission of the chassis ambient temperature trap.</p>
<p>cipSecTunnelStart (CISCO-IPSEC-FLOW-MONITOR-MIB)</p>	<p>cipSecTunLifeTime, cipSecTunLifeSize</p>	<p>The snmp-server enable traps ipsec start command is used to enable transmission of this trap.</p>
<p>cipSecTunnelStop (CISCO-IPSEC-FLOW-MONITOR-MIB)</p>	<p>cipSecTunActiveTime</p>	<p>The snmp-server enable traps ipsec stop command is used to enable transmission of this trap.</p>
<p>ciscoRasTooManySessions (CISCO-REMOTE-ACCESS-MONITOR-MIB)</p>	<p>crasNumSessions, crasNumUsers, crasMaxSessionsSupportable, crasMaxUsersSupportable, crasThrMaxSessions</p>	<p>The snmp-server enable traps remote-access session-threshold-exceeded command is used to enable transmission of these traps.</p>

Table 45-5 Supported Traps (Notifications) (continued)

clogMessageGenerated (CISCO-SYSLOG-MIB)	clogHistFacility, clogHistSeverity, clogHistMsgName, clogHistMsgText, clogHistTimestamp	Syslog messages are generated. The value of the clogMaxSeverity object is used to decide which syslog messages are sent as traps. The snmp-server enable traps syslog command is used to enable and disable transmission of these traps.
clrResourceLimitReached (CISCO-L4L7MODULE-RESOURCE-LIMIT-MIB)	clrResourceLimitValueType, clrResourceLimitMax, clogOriginIDType, clogOriginID	The snmp-server enable traps connection-limit-reached command is used to enable transmission of the connection-limit-reached notification. The clogOriginID object includes the context name from which the trap originated.
coldStart (SNMPv2-MIB)	—	The SNMP agent has started. The snmp-server enable traps snmp coldstart command is used to enable and disable transmission of these traps.
cpmCPURisingThreshold (CISCO-PROCESS-MIB)	cpmCPURisingThresholdValue, cpmCPUTotalMonIntervalValue, cpmCPUInterruptMonIntervalValue, cpmCPURisingThresholdPeriod, cpmProcessTimeCreated, cpmProcExtUtil5SecRev	The snmp-server enable traps cpu threshold rising command is used to enable transmission of the cpu threshold rising notification. The cpmCPURisingThresholdPeriod object is sent with the other objects.
entConfigChange (ENTITY-MIB)	—	The snmp-server enable traps entity config-change fru-insert fru-remove command is used to enable this notification. Note This notification is only sent in multimode when a security context is created or removed.
linkDown (IF-MIB)	ifIndex, ifAdminStatus, ifOperStatus	The linkdown trap for interfaces. The snmp-server enable traps snmp linkdown command is used to enable and disable transmission of these traps.
linkUp (IF-MIB)	ifIndex, ifAdminStatus, ifOperStatus	The linkup trap for interfaces. The snmp-server enable traps snmp linkup command is used to enable and disable transmission of these traps.

Table 45-5 Supported Traps (Notifications) (continued)

mteTriggerFired (DISMAN-EVENT-MIB)	mteHotTrigger, mteHotTargetName, mteHotContextName, mteHotOID, mteHotValue, cempMemPoolName, cempMemPoolHCUsed	The snmp-server enable traps memory-threshold command is used to enable the memory threshold notification. The mteHotOID is set to cempMemPoolHCUsed. The cempMemPoolName and cempMemPoolHCUsed objects are sent with the other objects.
mteTriggerFired (DISMAN-EVENT-MIB) Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.	mteHotTrigger, mteHotTargetName, mteHotContextName, mteHotOID, mteHotValue, ifHCInOctets, ifHCOutOctets, ifHighSpeed, entPhysicalName	The snmp-server enable traps interface-threshold command is used to enable the interface threshold notification. The entPhysicalName objects are sent with the other objects.
natPacketDiscard (NAT-MIB)	ifIndex	The snmp-server enable traps nat packet-discard command is used to enable the NAT packet discard notification. This notification is rate limited for 5 minutes and is generated when IP packets are discarded by NAT because mapping space is not available. The ifIndex gives the ID of the mapped interface.
warmStart (SNMPv2-MIB)	—	The snmp-server enable traps snmp warmstart command is used to enable and disable transmission of these traps.

SNMP Version 3

This section describes SNMP Version 3 and includes the following topics:

- [SNMP Version 3 Overview, page 45-15](#)
- [Security Models, page 45-16](#)
- [SNMP Groups, page 45-16](#)
- [SNMP Users, page 45-16](#)
- [SNMP Hosts, page 45-16](#)
- [Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software, page 45-16](#)

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, after you have added the **snmp-server host** command, 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 retrieval 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).

Configuring SNMP

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

- [Enabling SNMP, page 45-18](#)
- [Configuring SNMP Traps, page 45-20](#)
- [Configuring a CPU Usage Threshold, page 45-21](#)
- [Configuring a Physical Interface Threshold, page 45-21](#)
- [Using SNMP Version 1 or 2c, page 45-22](#)
- [Using SNMP Version 3, page 45-23](#)

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, enter the following command:

Command	Purpose
snmp-server enable Example: ciscoasa(config)# snmp-server enable	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 SNMP Traps” section on page 45-20.

Configuring SNMP Traps

To designate which traps that the SNMP agent generates and how they are collected and sent to NMSs, enter the following command:

Command	Purpose
<pre>snmp-server enable traps [all syslog snmp [authentication linkup linkdown coldstart warmstart] entity [config-change fru-insert fru-remove fan-failure cpu-temperature chassis-fan- failure power-supply-failure] chassis-temperature power-supply-presence power-supply-temperature] ikev2 [start stop] ipsec [start stop] remote-access [session-threshold-exceeded] connection-limit-reached cpu threshold rising interface-threshold memory-threshold nat [packet-discard]</pre>	<p>Sends individual traps, sets of traps, or all traps to the NMS. Enables syslog messages to be sent as traps to the NMS. The default configuration has all SNMP standard traps enabled, as shown in the example. To disable these traps, use the no snmp-server enable traps snmp command. If you enter this command and do not specify a trap type, the default is the syslog trap. By default, the syslog trap is enabled. The default SNMP traps continue to be enabled with the syslog trap. You need to configure both the logging history command and the snmp-server enable traps syslog command to generate traps from the syslog MIB. To restore the default enabling of SNMP traps, use the clear configure snmp-server command. All other traps are disabled by default.</p> <p>Keywords available in the admin context only:</p> <ul style="list-style-type: none"> • connection-limit-reached • entity • memory-threshold <p>Traps generated through the admin context only for physically connected interfaces in the system context:</p> <ul style="list-style-type: none"> • interface-threshold <p>Note The interface-threshold trap is not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.</p> <p>All other traps are available in the admin and user contexts in single mode. In multi-mode, the fan-failure trap, the power-supply-failure trap, and the cpu-temperature trap are generated only from the admin context, and not the user contexts (applies only to the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X). These traps do not apply to the ASA 5505.</p> <p>If the CPU usage is greater than the configured threshold value for the configured monitoring period, the cpu threshold rising trap is generated.</p> <p>When the used system context memory reaches 80 percent of the total system memory, the memory-threshold trap is generated from the admin context. For all other user contexts, this trap is generated when the used memory reaches 80 percent of the total system memory in that particular context.</p> <p>Note SNMP does not monitor voltage sensors.</p>

Example:

```
ciscoasa(config)# snmp-server enable
traps snmp authentication linkup
linkdown coldstart warmstart
```


What to Do Next

See the [“Configuring a CPU Usage Threshold”](#) section on page 45-21.

Configuring a CPU Usage Threshold

To configure the CPU usage threshold, enter the following command:

Command	Purpose
<pre>snmp cpu threshold rising threshold_value monitoring_period</pre> <p>Example: ciscoasa(config)# snmp cpu threshold rising 75% 30 minutes</p>	<p>Configures the threshold value for a high CPU threshold and the threshold monitoring period. To clear the threshold value and monitoring period of the CPU utilization, use the no form of this command. If the snmp cpu threshold rising command is not configured, the default for the high threshold level is over 70 percent, and the default for the critical threshold level is over 95 percent. The default monitoring period is set to 1 minute.</p> <p>You cannot configure the critical CPU threshold level, which is maintained at a constant 95 percent. Valid threshold values for a high CPU threshold range from 10 to 94 percent. Valid values for the monitoring period range from 1 to 60 minutes.</p>

What to Do Next

See the [“Configuring a Physical Interface Threshold”](#) section on page 45-21.

Configuring a Physical Interface Threshold

To configure the physical interface threshold, enter the following command:

Command	Purpose
<pre>snmp interface threshold threshold_value</pre> <p>Example: ciscoasa(config)# snmp interface threshold 75%</p> <p>Note Not supported on the ASA Services Module for Catalyst 6500 switches/7600 routers.</p>	<p>Configures the threshold value for an SNMP physical interface. To clear the threshold value for an SNMP physical interface, use the no form of this command. The threshold value is defined as a percentage of interface bandwidth utilization. Valid threshold values range from 30 to 99 percent. The default value is 70 percent.</p> <p>The snmp interface threshold command is available only in the admin context.</p> <p>Note Physical interface usage is monitored in single mode and multimode, and traps for physical interfaces in the system context are sent through the admin context. Only physical interfaces are used to compute threshold usage.</p>

What to Do Next

Choose one of the following:

- See the [“Using SNMP Version 1 or 2c”](#) section on page 45-22.
- See the [“Using SNMP Version 3”](#) section on page 45-23.

Using SNMP Version 1 or 2c

To configure parameters for SNMP Version 1 or 2c, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>snmp-server host interface) hostname ip_address} [trap poll] [community community-string] [version {1 2c username}] [udp-port port]</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server host mgmt 10.7.14.90 version 2</pre> <pre>ciscoasa(config)# snmp-server host corp 172.18.154.159 community public</pre>	<p>Specifies the recipient of an SNMP notification, indicates the interface from which traps are sent, and identifies the name and IP address of the NMS or SNMP manager that can connect to the ASA. The trap keyword limits the NMS to receiving traps only. The poll keyword limits the NMS to sending requests (polling) only. By default, SNMP traps are enabled. By default, the UDP port is 162. The community string is a shared secret key between the ASA or ASASM and the NMS. The key is a case-sensitive value up to 32 alphanumeric characters. Spaces are not permitted. The default community-string is public. The ASA uses this key to determine whether the incoming SNMP request is valid. For example, you could designate a site with a community string and then configure the ASA and the management station with the same string. The ASA and ASASM uses the specified string and does not respond to requests with an invalid community string. For more information about SNMP hosts, see the “SNMP Hosts” section on page 45-16.</p> <p>Note To receive traps, after you have added the snmp-server host command, make sure that you configure the user on the NMS with the same credentials as the credentials configured on the ASA and ASASM.</p>
Step 2	<pre>snmp-server community community-string</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server community onceuponatime</pre>	<p>Sets the community string, which is for use <i>only</i> with SNMP Version 1 or 2c.</p>
Step 3	<pre>snmp-server [contact location] text</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server location building 42</pre> <pre>ciscoasa(config)# snmp-server contact EmployeeA</pre>	<p>Sets the SNMP server location or contact information.</p>

What to Do Next

See the “Monitoring SNMP” section on page 45-26.

Using SNMP Version 3

To configure parameters for SNMP Version 3, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>snmp-server group group-name v3 [auth noauth priv]</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server group testgroup1 v3 auth</pre>	<p>Specifies a new SNMP group, which is for use <i>only</i> with SNMP Version 3. When a community string is configured, two additional groups with the name that matches the community string are autogenerated: one for the Version 1 security model and one for the Version 2 security model. For more information about security models, see the “Security Models” section on page 45-16. The auth keyword enables packet authentication. The noauth keyword indicates no packet authentication or encryption is being used. The priv keyword enables packet encryption and authentication. No default values exist for the auth or priv keywords.</p>
Step 2	<pre>snmp-server user username group-name {v3 [encrypted]} [auth {md5 sha}] auth-password [priv {des 3des aes} [128 192 256] priv-password</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server user testuser1 testgroup1 v3 auth md5 testpassword aes 128 mypassword</pre> <pre>ciscoasa(config)# snmp-server user testuser1 public v3 encrypted auth md5 00:11:22:33:44:55:66:77:88:99:AA: BB:CC:DD:EE:FF</pre>	<p>Configures a new user for an SNMP group, which is for use only with SNMP Version 3. The <i>username</i> argument is the name of the user on the host that belongs to the SNMP agent. The <i>group-name</i> argument is the name of the group to which the user belongs. The v3 keyword specifies that the SNMP Version 3 security model should be used and enables the use of the encrypted, priv, and the auth keywords. The encrypted keyword specifies the password in encrypted format. Encrypted passwords must be in hexadecimal format. The auth keyword specifies which authentication level (md5 or sha) should be used. The priv keyword specifies the encryption level. No default values for the auth or priv keywords, or default passwords exist. For the encryption algorithm, you can specify either the des, 3des, or aes keyword. You can also specify which version of the AES encryption algorithm to use: 128, 192, or 256. The <i>auth-password</i> argument specifies the authentication user password. The <i>priv-password</i> argument specifies the encryption user password.</p> <p>Note If you forget a password, you cannot recover it and you must reconfigure the user. You can specify a plain-text password or a localized digest. The localized digest must match the authentication algorithm selected for the user, which can be either MD5 or SHA. When the user configuration is displayed on the console or is written to a file (for example, the startup-configuration file), the localized authentication and privacy digests are always displayed instead of a plain-text password (see the second example). The minimum length for a password is 1 alphanumeric character; however, we recommend that you use at least 8 alphanumeric characters for security.</p>

	Command	Purpose
Step 3	<pre>snmp-server host interface {hostname ip_address} [trap poll] [community community-string] [version {1 2c 3 username}] [udp-port port]</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server host mgmt 10.7.14.90 version 3 testuser1</pre> <pre>ciscoasa(config)# snmp-server host mgmt 10.7.26.5 version 3 testuser2</pre>	<p>Specifies the recipient of an SNMP notification. Indicates the interface from which traps are sent. Identifies the name and IP address of the NMS or SNMP manager that can connect to the ASA. The trap keyword limits the NMS to receiving traps only. The poll keyword limits the NMS to sending requests (polling) only. By default, SNMP traps are enabled. By default, the UDP port is 162. The community string is a shared secret key between the ASA and the NMS. The key is a case-sensitive value up to 32 alphanumeric characters. Spaces are not permitted. The default community-string is public. The ASA and ASASM use this key to determine whether the incoming SNMP request is valid. For example, you could designate a site with a community string and then configure the ASA or ASASM and the NMS with the same string. The ASA and ASASM use the specified string and does not respond to requests with an invalid community string. For more information about SNMP hosts, see the “SNMP Hosts” section on page 45-16.</p> <p>Note When SNMP Version 3 hosts are configured on the ASA and ASASM, a user must be associated with that host. To receive traps, after you have added the snmp-server host command, make sure that you configure the user on the NMS with the same credentials as the credentials configured on the ASA and ASASM.</p>
Step 4	<pre>snmp-server [contact location] text</pre> <p>Example:</p> <pre>ciscoasa(config)# snmp-server location building 42</pre> <pre>ciscoasa(config)# snmp-server contact EmployeeA</pre>	<p>Sets the SNMP server location or contact information.</p>

What to Do Next

See the “Monitoring SNMP” section on page 45-26.

Troubleshooting Tips

To ensure that the SNMP process that receives incoming packets from the NMS is running, enter the following command:

```
ciscoasa(config)# show process | grep snmp
```

To capture syslog messages from SNMP and have them appear on the ASA or ASASM console, enter the following commands:

```
ciscoasa(config)# logging list snmp message 212001-212015
ciscoasa(config)# logging console snmp
```

To make sure that the SNMP process is sending and receiving packets, enter the following commands:

```
ciscoasa(config)# clear snmp-server statistics
ciscoasa(config)# show snmp-server statistics
```

The output is based on the SNMP group of the SNMPv2-MIB.

To make sure that SNMP packets are going through the ASA or ASASM and to the SNMP process, enter the following commands:

```
ciscoasa(config)# clear asp drop
ciscoasa(config)# show asp drop
```

If the NMS cannot request objects successfully or is not handing incoming traps from the ASA or ASASM correctly, use a packet capture to isolate the problem, by entering the following commands:

```
hostname (config)# access-list snmp permit udp any eq snmptrap any
hostname (config)# access-list snmp permit udp any any eq snmp
hostname (config)# capture snmp type raw-data access-list snmp interface mgmt
hostname (config)# copy /pcap capture:snmp tftp://192.0.2.5/exampledir/snmp.pcap
```

If the ASA or ASASM is not performing as expected, obtain information about network topology and traffic by doing the following:

- For the NMS configuration, obtain the following information:
 - Number of timeouts
 - Retry count
 - Engine ID caching
 - Username and password used
- Run the following commands:
 - **show block**
 - **show interface**
 - **show process**
 - **show cpu**

If a fatal error occurs, to help in reproducing the error, send a traceback file and the output of the **show tech-support** command to Cisco TAC.

If SNMP traffic is not being allowed through the ASA or ASASM interfaces, you might also need to permit ICMP traffic from the remote SNMP server using the **icmp permit** command.

For the ASA 5580, differences may appear in the physical interface statistics output and the logical interface statistics output between the **show interface** command and the **show traffic** command.

Interface Types and Examples

The interface types that produce SNMP traffic statistics include the following:

- Logical—Statistics collected by the software driver, which are a subset of physical statistics.
- Physical—Statistics collected by the hardware driver. Each physical named interface has a set of logical and physical statistics associated with it. Each physical interface may have more than one VLAN interface associated with it. VLAN interfaces only have logical statistics.



Note For a physical interface that has multiple VLAN interfaces associated with it, be aware that SNMP counters for ifInOctets and ifOutOctets OIDs match the aggregate traffic counters for that physical interface.

- VLAN-only—SNMP uses logical statistics for ifInOctets and ifOutOctets.

The examples in [Table 45-6](#) show the differences in SNMP traffic statistics. Example 1 shows the difference in physical and logical output statistics for the **show interface** command and the **show traffic** command. Example 2 shows output statistics for a VLAN-only interface for the **show interface** command and the **show traffic** command. The example shows that the statistics are close to the output that appears for the **show traffic** command.

Table 45-6 SNMP Traffic Statistics for Physical and VLAN Interfaces

Example 1	Example 2
<pre> ciscoasa# show interface GigabitEthernet3/2 interface GigabitEthernet3/2 description fullt-mgmt nameif mgmt security-level 10 ip address 10.7.14.201 255.255.255.0 management-only ciscoasa# show traffic (Condensed output) Physical Statistics GigabitEthernet3/2: received (in 121.760 secs) 36 packets 3428 bytes 0 pkts/sec 28 bytes/sec Logical Statistics mgmt: received (in 117.780 secs) 36 packets 2780 bytes 0 pkts/sec 23 bytes/sec The following examples show the SNMP output statistics for the management interface and the physical interface. The ifInOctets value is close to the physical statistics output that appears in the show traffic command output but not to the logical statistics output. ifIndex of the mgmt interface: IF-MIB::ifDescr.6 = Adaptive Security Appliance 'mgmt' interface ifInOctets that corresponds to the physical interface statistics: IF-MIB::ifInOctets.6 = Counter32:3246 </pre>	<pre> ciscoasa# show interface GigabitEthernet0/0.100 interface GigabitEthernet0/0.100 vlan 100 nameif inside security-level 100 ip address 10.7.1.101 255.255.255.0 standby 10.7.1.102 ciscoasa# show traffic inside received (in 9921.450 secs) 1977 packets 126528 bytes 0 pkts/sec 12 bytes/sec transmitted (in 9921.450 secs) 1978 packets 126556 bytes 0 pkts/sec 12 bytes/sec ifIndex of VLAN inside: IF-MIB::ifDescr.9 = Adaptive Security Appliance 'inside' interface IF-MIB::ifInOctets.9 = Counter32: 126318 </pre>

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 45-27](#)
- [SNMP Monitoring, page 45-27](#)

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, enter one of the following commands:

Command	Purpose
<code>show running-config [default] snmp-server</code>	Shows all SNMP server configuration information.
<code>show running-config snmp-server group</code>	Shows SNMP group configuration settings.
<code>show running-config snmp-server host</code>	Shows configuration settings used by SNMP to control messages and notifications sent to remote hosts.
<code>show running-config snmp-server user</code>	Shows SNMP user-based configuration settings.
<code>show snmp-server engineid</code>	Shows the ID of the SNMP engine configured.
<code>show snmp-server group</code>	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.
<code>show snmp-server statistics</code>	Shows the configured characteristics of the SNMP server. To reset all SNMP counters to zero, use the clear snmp-server statistics command.
<code>show snmp-server user</code>	Shows the configured characteristics of users.

Examples

The following example shows how to display SNMP server statistics:

```
ciscoasa(config)# show snmp-server statistics
0 SNMP packets input
  0 Bad SNMP version errors
  0 Unknown community name
  0 Illegal operation for community name supplied
  0 Encoding errors
  0 Number of requested variables
  0 Number of altered variables
  0 Get-request PDUs
  0 Get-next PDUs
  0 Get-bulk PDUs
```

```

    0 Set-request PDUs (Not supported)
0 SNMP packets output
    0 Too big errors (Maximum packet size 512)
    0 No such name errors
    0 Bad values errors
    0 General errors
    0 Response PDUs
    0 Trap PDUs

```

The following example shows how to display the SNMP server running configuration:

```

ciscoasa(config)# show running-config snmp-server
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart

```

Configuration Examples for SNMP

This section includes the following topics:

- [Configuration Example for SNMP Versions 1 and 2c, page 45-28](#)
- [Configuration Example for SNMP Version 3, page 45-28](#)

Configuration Example for SNMP Versions 1 and 2c

The following example shows how the ASA can receive SNMP requests from host 192.0.2.5 on the inside interface but does not send any SNMP syslog requests to any host:

```

ciscoasa(config)# snmp-server host 192.0.2.5
ciscoasa(config)# snmp-server location building 42
ciscoasa(config)# snmp-server contact EmployeeA
ciscoasa(config)# snmp-server community ohwhatakeyisthee

```

Configuration Example for SNMP Version 3

The following example shows how the ASA can receive SNMP requests using the SNMP Version 3 security model, which requires that the configuration follow this specific order: group, followed by user, followed by host:

```

ciscoasa(config)# snmp-server group v3 vpn-group priv
ciscoasa(config)# snmp-server user admin vpn group v3 auth sha letmein priv 3des cisco123
ciscoasa(config)# snmp-server host mgmt 10.0.0.1 version 3 priv admin

```


Where to Go Next

To configure the syslog server, see [Chapter 44, “Configuring Logging.”](#)

Additional References

For additional information related to implementing SNMP, see the following sections:

- [RFCs for SNMP Version 3, page 45-29](#)
- [MIBs, page 45-29](#)
- [Application Services and Third-Party Tools, page 45-31](#)

RFCs for SNMP Version 3

RFC	Title
3410	<i>Introduction and Applicability Statements for Internet Standard Management Framework</i>
3411	<i>An Architecture for Describing SNMP Management Frameworks</i>
3412	<i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>
3413	<i>Simple Network Management Protocol (SNMP) Applications</i>
3414	<i>User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMP)</i>
3826	<i>The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model</i>

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, enter the following command:

```
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.	snmpInTooBig
[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.	snmpInReadOnly
[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.	snmpOutTooBig
[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 45-7 lists each feature change and the platform release in which it was implemented.

Table 45-7 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.
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 introduced or modified the following commands: show snmp-server engineid , show snmp-server group , show snmp-server user , snmp-server group , snmp-server user , snmp-server host .
Password encryption	8.3(1)	Supports password encryption. We modified the following commands: snmp-server community , snmp-server host .

Table 45-7 Feature History for SNMP (continued)

Feature Name	Platform Releases	Feature Information
SNMP traps and MIBs	8.4(1)	<p>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.</p> <p>The entPhysicalTable reports entries for sensors, fans, power supplies, and related components.</p> <p>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.</p> <p>Supports the following additional traps: ceSensorExtThresholdNotification, clrResourceLimitReached, cpmCPURisingThreshold, mteTriggerFired, natPacketDiscard, warmStart.</p> <p>We introduced or modified the following commands: snmp cpu threshold rising, snmp interface threshold, snmp-server enable traps.</p>
IF-MIB ifAlias OID support	8.2(5)/8.4(2)	The ASA now supports the ifAlias OID. When you browse the IF-MIB, the ifAlias OID will be set to the value that has been set for the interface description.
ASA Services Module (ASASM)	8.5(1)	<p>The ASASM supports all MIBs and traps that are present in 8.4(1), except for the following:</p> <p>Unsupported MIBs in 8.5(1):</p> <ul style="list-style-type: none"> • 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). <p>Unsupported traps in 8.5(1):</p> <ul style="list-style-type: none"> • 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)	<p>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.</p> <p>We modified the following command: snmp-server enable traps.</p>

Table 45-7 Feature History for SNMP (continued)

Feature Name	Platform Releases	Feature Information
VPN-related MIBs	9.0(1)	<p>An updated version of the CISCO-IPSEC-FLOW-MONITOR-MIB.my MIB has been implemented to support the next generation encryption feature.</p> <p>The following MIBs have been enabled for the ASASM:</p> <ul style="list-style-type: none"> • 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.



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 46-1](#)
- [Licensing Requirements for NSEL, page 46-4](#)
- [Prerequisites for NSEL, page 46-4](#)
- [Guidelines and Limitations, page 46-4](#)
- [Configuring NSEL, page 46-5](#)
- [Monitoring NSEL, page 46-10](#)
- [Configuration Examples for NSEL, page 46-12](#)
- [Where to Go Next, page 46-13](#)
- [Additional References, page 46-13](#)
- [Feature History for NSEL, page 46-14](#)

Information About NSEL

This section includes the following topics:

- [Using NSEL and Syslog Messages, page 46-2](#)
- [Using NSEL in Clustering, page 46-3](#)

The ASA and ASASM support NetFlow Version 9 services. For more information about NetFlow services, see the “[RFCs](#)” section on page 46-14.

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 46-1](#) 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 46-1](#) redundant. In the interest of performance, we recommend that you disable redundant syslog messages, because the same information is exported through NetFlow. You can enable or disable individual syslog messages by following the procedure in the [“Disabling and Reenabling NetFlow-related Syslog Messages”](#) section on page 46-9.

Table 46-1 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 8, “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:

- [Configuring NSEL Collectors, page 46-5](#)
- [Configuring Flow-Export Actions Through Modular Policy Framework, page 46-6](#)
- [Configuring Template Timeout Intervals, page 46-7](#)
- [Changing the Time Interval for Sending Flow-Update Events to a Collector, page 46-8](#)
- [Delaying Flow-Create Events, page 46-9](#)
- [Disabling and Reenabling NetFlow-related Syslog Messages, page 46-9](#)
- [Clearing Runtime Counters, page 46-10](#)

Configuring NSEL Collectors

To configure NSEL collectors, enter the following command:

Command	Purpose
<p>flow-export destination <i>interface-name</i> <i>ipv4-address hostname udp-port</i></p> <p>Example: hostname (config)# flow-export destination inside 209.165.200.225 2002</p>	<p>Adds, edits, or deletes an NSEL collector to which NetFlow packets are sent. The destination keyword indicates that a NSEL collector is being configured. The <i>interface-name</i> argument is the name of the ASA and ASA Services Module interface through which the collector is reached. The <i>ipv4-address</i> argument is the IP address of the machine running the collector application. The <i>hostname</i> argument is the destination IP address or name of the collector. The <i>udp-port</i> argument is the UDP port number to which NetFlow packets are sent. You can configure a maximum of five collectors. After a collector is configured, template records are automatically sent to all configured NSEL collectors.</p> <p>Note Make sure that collector applications use the Event Time field to correlate events.</p>

What to Do Next

See the “[Configuring Flow-Export Actions Through Modular Policy Framework](#)” section on page 46-6.

Configuring Flow-Export Actions Through Modular Policy Framework

To export NSEL events by defining all classes with flow-export actions, perform the following steps:

	Command	Purpose
Step 1	<p>class-map <i>flow_export_class</i></p> <p>Example: hostname (config-pmap)# class-map flow_export_class</p>	Defines the class map that identifies traffic for which NSEL events need to be exported. The <i>flow_export_class</i> argument is the name of the class map.
Step 2	<p>Choose one of the following options:</p> <p>match access-list <i>flow_export_acl</i></p> <p>Example: hostname (config-cmap)# match access-list flow_export_acl</p> <p>match any</p> <p>Example: hostname (config-cmap)# match any</p>	<p>Configures the ACL to match specific traffic. The <i>flow_export_acl</i> argument is the name of the ACL.</p> <p>Matches any traffic.</p>
Step 3	<p>policy-map <i>flow_export_policy</i></p> <p>Example: hostname(config)# policy-map flow_export_policy</p>	<p>Defines the policy map to apply flow-export actions to the defined classes. The <i>flow_export_policy</i> argument is the name of the policy map.</p> <p>If you create a new policy map and apply it globally according to Step 6, the remaining inspection policies are deactivated.</p> <p>Alternatively, to insert a NetFlow class in the existing policy, enter the class flow_export_class command after the policy-map global_policy command.</p> <p>For more information about creating or modifying the Modular Policy Framework, see Chapter 1, “Configuring a Service Policy Using the Modular Policy Framework,” in the firewall configuration guide.</p>
Step 4	<p>class <i>flow_export_class</i></p> <p>Example: hostname (config-pmap)# class flow_export_class</p>	Defines the class to apply flow-export actions. The <i>flow_export_class</i> argument is the name of the class.

	Command	Purpose
Step 5	<pre>flow-export event-type event-type destination flow_export_host1 [flow_export_host2]</pre> <p>Example: hostname (config-pmap-c)# flow-export event-type all destination 209.165.200.230 </p>	Configures a flow-export action. The event_type keyword is the name of the supported event being filtered. The destination keyword is the IP address of the configured collector. The <i>flow_export_host</i> argument is the IP address of a host.
Step 6	<pre>service-policy flow_export_policy global</pre> <p>Example: hostname (config)# service-policy flow_export_policy global </p>	Adds or edits the service policy globally. The <i>flow_export_policy</i> argument is the name of the policy map.

What to Do Next

See the [“Configuring Template Timeout Intervals”](#) section on page 46-7.

Configuring Template Timeout Intervals

To configure template timeout intervals, enter the following command:

Command	Purpose
<pre>flow-export template timeout-rate minutes</pre> <p>Example: hostname (config)# flow-export template timeout-rate 15 </p>	Specifies the interval at which template records are sent to all configured output destinations. The template keyword indicates the template-specific configurations. The timeout-rate keyword specifies the time before templates are resent. The <i>minutes</i> argument specifies the time interval in minutes at which the templates are resent. The default value is 30 minutes.

What to Do Next

See the [“Changing the Time Interval for Sending Flow-Update Events to a Collector”](#) section on page 46-8.

Changing the Time Interval for Sending Flow-Update Events to a Collector

To change the time interval at which periodic flow-update events are to be sent to a collector, enter the following command:

Command	Purpose
<p>flow-export active refresh-interval <i>value</i></p> <p>Example: hostname (config)# flow-export active refresh-interval 30</p>	<p>Configures NetFlow parameters for active connections. The <i>value</i> argument specifies the time interval between flow-update events in minutes. Valid values are from 1 - 60 minutes. The default value is 1 minute.</p> <p>If you have already configured the flow-export delay flow-create command, and you then configure the flow-export active refresh-interval command with an interval value that is not at least 5 seconds more than the delay value, the following warning message appears at the console:</p> <pre>WARNING: The current delay flow-create value configuration may cause flow-update events to appear before flow-creation events.</pre> <p>If you have already configured the flow-export active refresh-interval command, and you then configure the flow-export delay flow-create command with a delay value that is not at least 5 seconds less than the interval value, the following warning message appears at the console:</p> <pre>WARNING: The current delay flow-create value configuration may cause flow-update events to appear before flow-creation events.</pre>

What to Do Next

See the [“Delaying Flow-Create Events”](#) section on page 46-9.

Delaying Flow-Create Events

To delay the sending of flow-create events, enter the following command:

Command	Purpose
flow-export delay flow-create <i>seconds</i> Example: hostname (config)# flow-export delay flow-create 10	Delays the sending of a flow-create event by the specified number of seconds. The <i>seconds</i> argument indicates the amount of time allowed for the delay in seconds. If this command is not configured, there is no delay, and the flow-create event is exported as soon as the flow is created. If the flow is torn down before the configured delay, the flow-create event is not sent; an extended flow teardown event is sent instead.

What to Do Next

See the [“Disabling and Reenabling NetFlow-related Syslog Messages”](#) section on page 46-9.

Disabling and Reenabling NetFlow-related Syslog Messages

To disable and reenablenet NetFlow-related syslog messages, perform the following steps:

	Command	Purpose
Step 1	logging flow-export-syslogs disable Example: hostname(config)# logging flow-export-syslogs disable	Disables syslog messages that have become redundant because of NSEL. Note Although you execute this command in global configuration mode, it is not stored in the configuration. Only the no logging message xxxxxx commands are stored in the configuration.
Step 2	logging message xxxxxx Example: hostname(config)# logging message 302013	Reenables syslog messages individually, where <i>xxxxxx</i> is the specified syslog message that you want to reenablenet.
Step 3	logging flow-export-syslogs enable Example: hostname(config)# logging flow-export-syslogs enable	Reenables all NSEL events at the same time.

What to Do Next

See the [“Clearing Runtime Counters”](#) section on page 46-10.

Clearing Runtime Counters

To reset runtime counters, enter the following command:

Command	Purpose
<code>clear flow-export counters</code>	Resets all runtime counters for NSEL to zero.
Example: <code>hostname# clear flow-export counters</code>	

What to Do Next

See the [“Monitoring NSEL”](#) section on page 46-10.

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 46-2.

NSEL Monitoring Commands

To monitor NSEL, enter one of the following commands:

Command	Purpose
<code>show flow-export counters</code>	Shows runtime counters, including statistical data and error data, for NSEL.
<code>show logging flow-export-syslogs</code>	Lists all syslog messages that are captured by NSEL events.
<code>show running-config flow-export</code>	Shows the currently configured NetFlow commands.
<code>show running-config logging</code>	Shows disabled syslog messages, which are redundant syslog messages, because they export the same information through NetFlow.

Examples

The following example shows how to display flow-export counters:

```
hostname (config)# show flow-export counters
```



```
destination: inside 209.165.200.225 2055
```

```
Statistics:
  packets sent           250
Errors:
  block allocation errors      0
  invalid interface           0
  template send failure       0
  no route to collector       0
  source port allocation      0
```

The following example shows how to display the flow-export active configuration:

```
ciscoasa (config)# show running-config flow-export active
flow-export active refresh-interval 2
```

The following example shows how to display the flow-export delay configuration:

```
hostname (config)# show running-config flow-export delay
flow-export delay flow-create 30
```

The following example shows how to display the flow-export destination configurations:

```
hostname (config)# show running-config flow-export destination
flow-export destination inside 192.68.10.70 9996
```

The following example shows how to display the flow-export template configuration:

```
hostname (config)# show running-config flow-export template
flow-export template timeout-rate 1
```

The following example shows how to display flow-export syslog messages:

```
hostname# show logging flow-export-syslogs
```

Syslog ID	Type	Status
302013	Flow Created	Enabled
302015	Flow Created	Enabled
302017	Flow Created	Enabled
302020	Flow Created	Enabled
302014	Flow Deleted	Enabled
302016	Flow Deleted	Enabled
302018	Flow Deleted	Enabled
302021	Flow Deleted	Enabled
106015	Flow Denied	Enabled
106023	Flow Denied	Enabled
313001	Flow Denied	Enabled
313008	Flow Denied	Enabled
710003	Flow Denied	Enabled
106100	Flow Created/Denied	Enabled

The following example shows how to display current syslog message settings:

```
hostname (config)# show running-config logging
```

```
no logging message 313008
no logging message 313001
```

Configuration Examples for NSEL

The following examples show how to filter NSEL events, with the specified collectors already configured:

- **flow-export destination inside 209.165.200.2055**
- **flow-export destination outside 209.165.201.29 2055**
- **flow-export destination outside 209.165.201.27 2055**

Log all events between hosts 209.165.200.224 and hosts 209.165.201.224 to 209.165.200.230, and log all other events to 209.165.201.29:

```
hostname (config)# access-list flow_export_acl permit ip host 209.165.200.224 host
209.165.201.224
hostname (config)# class-map flow_export_class
hostname (config-cmap)# match access-list flow_export_acl
hostname (config)# policy-map flow_export_policy
hostname (config-pmap)# class flow_export_class
hostname (config-pmap-c)# flow-export event-type all destination 209.165.200.230
hostname (config-pmap)# class class-default
hostname (config-pmap-c)# flow-export event-type all destination 209.165.201.29
hostname (config)# service-policy flow_export_policy global
```

Log flow-create events to 209.165.200.230, flow-teardown events to 209.165.201.29, flow-denied events to 209.165.201.27, and flow-update events to 209.165.200.230:

```
hostname (config)# policy-map flow_export_policy
hostname (config-pmap)# class class-default
hostname (config-pmap-c)# flow-export event-type flow-creation destination 209.165.200.230
hostname (config-pmap-c)# flow-export event-type flow-teardown destination 209.165.201.29
hostname (config-pmap-c)# flow-export event-type flow-denied destination 209.165.201.27
hostname (config-pmap-c)# flow-export event-type flow-update destination 209.165.200.230
hostname (config)# service-policy flow_export_policy global
```

Log flow-create events between hosts 209.165.200.224 and 209.165.200.230 to 209.165.201.29, and log all flow-denied events to 209.165.201.27:

```
hostname (config)# access-list flow_export_acl permit ip host 209.165.200.224 host
209.165.200.230
hostname (config)# class-map flow_export_class
hostname (config)# match access-list flow_export_acl
hostname (config)# policy-map flow_export_policy
hostname (config-pmap)# class flow_export_class
hostname (config-pmap-c)# flow-export event-type flow-creation destination 209.165.200.29
hostname (config-pmap-c)# flow-export event-type flow-denied destination 209.165.201.27
hostname (config-pmap)# class class-default
hostname (config-pmap-c)# flow-export event-type flow-denied destination 209.165.201.27
hostname (config)# service-policy flow_export_policy global
```



Note

You must enter the following command:

```
hostname (config-pmap-c)# flow-export event-type flow-denied destination 209.165.201.27
```

for *flow_export_acl*, because traffic is not checked after the first match, and you must explicitly define the action to log flow-denied events that match *flow_export_acl*.

Log all traffic except traffic between hosts 209.165.201.27 and 209.165.201.50 to 209.165.201.27:

```
hostname (config)# access-list flow_export_acl deny ip host 209.165.201.27 host
209.165.201.50
```

```
hostname (config)# access-list flow_export_acl permit ip any any
hostname (config)# class-map flow_export_class
hostname (config-cmap)# match access-list flow_export_acl
hostname (config)# policy-map flow_export_policy
hostname (config-pmap)# class flow_export_class
hostname (config-pmap-c)# flow-export event-type all destination 209.165.201.27
hostname (config)# service-policy flow_export_policy global
```

Where to Go Next

To configure the syslog server, see [Chapter 44, “Configuring Logging.”](#)

Additional References

For additional information related to implementing NSEL, see the following sections:

- [Related Documents, page 46-14](#)
- [RFCs, page 46-14](#)

Related Documents

Related Topic	Document Title
Using NSEL and Syslog Messages, page 46-2	<i>syslog messages guide</i>
Information about the implementation of NSEL on the ASA and ASA Services Module	<i>Cisco ASA 5500 Series Implementation Note for NetFlow Collectors</i> 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 46-2](#) lists each feature change and the platform release in which it was implemented..

Table 46-2 Feature History for NSEL

Feature Name	Platform Releases	Feature Information
NetFlow	8.1(1)	<p>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).</p> <p>We introduced the following commands: clear flow-export counters, flow-export enable, flow-export destination, flow-export template timeout-rate, logging flow-export syslogs enable, logging flow-export syslogs disable, show flow-export counters, show logging flow-export-syslogs.</p>
NetFlow Filtering	8.1(2)	<p>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.</p> <p>We modified the following commands: class, class-map, flow-export event-type destination, match access-list, policy-map, service-policy.</p> <p>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.</p> <p>We introduced the following command: flow-export delay flow-create.</p>
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		<p>A new NetFlow error counter, source port allocation failure, has been added.</p> <p>We modified the following command: show flow-export counters.</p> <p>Note The flow-update event feature is not available in Version 9.0(1).</p>
NSEL	9.1(2)	<p>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.</p> <p>We introduced the following command: flow-export active refresh-interval.</p> <p>We modified the following command: flow-export event-type.</p>



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 47-1](#)
- [Licensing Requirements for Anonymous Reporting and Smart Call Home, page 47-4](#)
- [Prerequisites for Smart Call Home and Anonymous Reporting, page 47-4](#)
- [Guidelines and Limitations, page 47-4](#)
- [Configuring Anonymous Reporting and Smart Call Home, page 47-5](#)
- [Monitoring Anonymous Reporting and Smart Call Home, page 47-22](#)
- [Configuration Example for Smart Call Home, page 47-23](#)
- [Feature History for Anonymous Reporting and Smart Call Home, page 47-24](#)

Information About Anonymous Reporting and Smart Call Home

This section includes the following topics:

- [Information About Anonymous Reporting, page 47-2](#)
- [Information About Smart Call Home, page 47-3](#)

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 47-6](#) or the [“Configuring Smart Call Home” section on page 47-6](#).

Information About Smart Call Home

When fully configured, Smart Call Home detects issues at your site and reports them back to Cisco or through other user-defined channels (such as e-mail or directly to you), often before you know that these issues exist. Depending upon the seriousness of these problems, Cisco responds to 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 47-2](#) and the [“Configuring the DNS Server” section on page 13-13](#).

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 47-6](#)
- [Configuring Smart Call Home, page 47-6](#)

Configuring Anonymous Reporting

To configure Anonymous Reporting and securely provide minimal error and health information to Cisco, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home reporting anonymous Example: ciscoasa(config)# call-home reporting anonymous	Enables the Anonymous Reporting feature and creates a new anonymous profile. Entering this command creates a trust point and installs a certificate that is used to verify the identity of the Cisco web server.
Step 2	call-home test reporting anonymous Example: ciscoasa(config)# call-home test reporting anonymous INFO: Sending test message to https://tools.cisco.com/its/service/oddce/ services/DDCEService... INFO: Succeeded	(Optional) Ensures that you have connectivity to the server and that your system can send messages. A success or error message returns test results.

Configuring Smart Call Home

This section includes the following topics:

- [Enabling Smart Call Home, page 47-6](#)
- [Declaring and Authenticating a CA Trust Point, page 47-7](#)
- [Subscribing to Alert Groups, page 47-8](#)
- [Optional Configuration Procedures, page 47-15](#)

Enabling Smart Call Home

To enable Smart Call Home and activate your call-home profile, perform the following steps:

Step 1	service call-home Example: hostname(config)# service call-home	Enables the Smart Call Home service.
Step 2	call-home Example: hostname(config)# call-home	Enters call-home configuration mode.

Step 3 Example: <pre>hostname(cfg-call-home)# contact-email-addr username@example.com</pre>	<pre>contact-email-addr email</pre>	Configures the mandatory contact address. The address should be the Cisco.com ID account associated with the device. This account is the e-mail address that you used to register the ASA with Cisco on Cisco.com.
Step 4 Example: <pre>hostname(cfg-call-home)# profile CiscoTAC-1</pre>	<pre>profile profile-name</pre>	Enables the profile. The default profile name is CiscoTAC-1.
Step 5 Example: <pre>hostname(cfg-call-home-profile)# active</pre>	<pre>active</pre>	Activates the call home profile. To disable this profile, enter the no active command.
Step 6 Example: <pre>hostname(cfg-call-home-profile)# destination transport-method http</pre>	<pre>destination transport-method http</pre>	Configures the destination transport method for the smart call-home message receiver. The default destination transport method is e-mail. To configure e-mail, see the “Enabling Smart Call Home” section on page 47-6 .

Declaring and Authenticating a CA Trust Point

If Smart Call Home is configured to send messages to a web server through HTTPS, you need to configure the ASA to trust the certificate of the web server or the certificate of the Certificate Authority (CA) that issued the certificate. The Cisco Smart Call Home Production server certificate is issued by Verisign. The Cisco Smart Call Home Staging server certificate is issued by the Digital Signature Trust Co.

Detailed Steps

To declare and authenticate the Cisco server security certificate and establish communication with the Cisco HTTPS server for Smart Call Home service, perform the following steps:

Step 1 Example: <pre>ciscoasa(config)# changeto context contextA</pre>	<pre>changeto context admincontext</pre>	(Multiple Context Mode only) Installs the certificate in the admin context.
Step 2 Example: <pre>ciscoasa(config)# crypto ca trustpoint cisco</pre>	<pre>crypto ca trustpoint trustpoint-name</pre>	Configures a trust point and prepares for certificate enrollment. Note If you use HTTP as the transport method, you must install a security certificate through a trust point, which is required for HTTPS. Find the specific certificate to install at the following URL: http://www.cisco.com/en/US/docs/switches/lan/smart_call_home/SCH31_Ch6.html#wp1035380

Step 3	enroll terminal Example: ciscoasa(ca-trustpoint)# enroll terminal	Specifies a manual cut-and-paste method of certificate enrollment.
Step 4	crypto ca authenticate trustpoint Example: ciscoasa(ca-trustpoint)# crypto ca authenticate cisco	Authenticates the named CA. The CA name should match the trust point name specified in the crypto ca trustpoint command. At the prompt, paste the security certificate text.
Step 5	quit Example: ciscoasa(ca-trustpoint)# quit %Do you accept this certificate [yes/no]: yes	Specifies the end of the security certificate text and confirms acceptance of the entered security certificate.

Subscribing to Alert Groups

An alert group is a predefined subset of the Smart Call Home alerts that are supported on the ASA. Different types of Smart Call Home alerts are grouped into different alert groups, depending on their type. Each alert group reports the output of certain CLIs. The supported Smart Call Home alert groups are the following:

- syslog
- diagnostic
- environment
- inventory
- configuration
- threat
- snapshot
- telemetry
- test

This section includes the following topics:

- [Attributes of Alert Groups, page 47-9](#)
- [Information Sent to Cisco by Alert Groups, page 47-9](#)
- [Information About the Message Severity Threshold, page 47-11](#)
- [Information About Subscription Profiles, page 47-11](#)
- [Configuring the Environment and Snapshot Alert Groups, page 47-12](#)
- [Configuring Alert Group Subscription, page 47-13](#)

Attributes of Alert Groups

Alert groups have the following attributes:

- Events first register with one alert group.
- A group can associate with multiple events.
- You can subscribe to specific alert groups.
- You can enable and disable specific alert groups. The default setting is enabled for all alert groups.
- The diagnostic and environment alert groups support subscription for periodic messages.
- The syslog alert group supports message ID-based subscription.
- You can configure a threshold for CPU and memory usage for the environment alert group. When a certain parameter has exceeded a predefined threshold, a message is sent. Most of the threshold values are platform-dependent and cannot be changed.
- You configure the snapshot alert group to send the output of CLIs that you specify.

Information Sent to Cisco by Alert Groups

Messages are sent to Cisco periodically and whenever the ASA reloads. These messages are categorized by alert groups.

Inventory alerts consist of output from the following commands:

- **show version**—Displays the ASA software version, hardware configuration, license key, and related uptime data for the device.
- **show inventory**—Retrieves and displays inventory information about each Cisco product that is installed in the networking device. Each product is identified by unique device information, called the UDI, which is a combination of three separate data elements: the product identifier (PID), the version identifier (VID), and the serial number (SN).
- **show failover state**—Displays the failover state of both units in a failover pair. The information displayed includes the primary or secondary status of the unit, the Active/Standby status of the unit, and the last reported reason for failover.
- **show module**—Shows information about any modules installed on the ASAs, for example, information about an AIP SSC installed on the ASA 5505 or information about an SSP installed on the ASA 5585-X, and information about an IPS SSP installed on an ASA 5585-X.
- **show environment**—Shows system environment information for ASA system components, such as hardware operational status for the chassis, drivers, fans, and power supplies, as well as temperature status, voltage, and CPU usage.

Configuration alerts consist of output from the following commands:

- **show context**—Shows allocated interfaces and the configuration file URL, the number of contexts configured, or if you enable Anonymous Reporting in the system execution space, from a list of all contexts.
- **show call-home registered-module status**—Shows the registered module status. If you use system configuration mode, the command displays system module status based on the entire device, not per context.
- **show running-config**—Shows the configuration that is currently running on the ASA.
- **show startup-config**—Show the startup configuration.
- **show access-list | include elements**—Shows the hit counters and a timestamp value for an access list.

Diagnostic alerts consist of output from the following commands:

- **show failover**—Displays information about the failover status of the unit.
- **show interface**—Displays interface statistics.
- **show cluster info**—Displays cluster information.
- **show cluster history**—Displays the cluster history.
- **show crashinfo** (truncated)—After an unexpected software reload, the device sends a modified crash information file with only the traceback section of the file included, so only function calls, register values, and stack dumps are reported to Cisco.
- **show tech-support no-config**—Displays the information that is used for diagnosis by technical support analysts.

Environment alerts consist of output from the following command:

- **show environment**—Shows system environment information for ASA system components, such as hardware operational status for the chassis, drivers, fans, and power supplies, as well as temperature status, voltage, and CPU usage.
- **show cpu usage**—Displays CPU usage information.
- **show memory detail**—Displays details of the free and allocated system memory.

Threat alerts consist of output from the following commands:

- **show threat-detection rate**—Displays threat detection statistics.
- **show threat-detection shun**—Displays currently shunned hosts.
- **show shun**—Displays shun information.
- **show dynamic-filter reports top**—Generates reports of the top 10 malware sites, ports, and infected hosts classified by the Botnet Traffic Filter.

Snapshot alerts may consist of output from the following commands (for example):

- **show conn count**—Shows the number of active connections.
- **show asp drop**—Shows the accelerated security path dropped packets or connections.

Telemetry alerts consist of output from the following commands:

- **show perfmon detail**—Shows ASA performance details.
- **show traffic**—Displays interface transmit and receive activity.
- **show conn count**—Shows the number of active connections.
- **show vpn-sessiondb summary**—Shows VPN session summary information.
- **show vpn load-balancing**—Displays the runtime statistics for the VPN load-balancing virtual cluster configuration.
- **show local-host | include interface**—Shows the network states of local hosts.
- **show memory**—Displays a summary of the maximum physical memory and current free memory available to the operating system.
- **show context**—Shows allocated interfaces and the configuration file URL, the number of contexts configured, or if you enable Anonymous Reporting in the system execution space, from a list of all contexts.
- **show access-list | include elements**—Shows the hit counters and a timestamp value for an access list.
- **show interface**—Displays interface statistics.

- **show threat-detection statistics protocol**—Shows IP protocol statistics.
- **show phone-proxy media-sessions count**—Displays the number of corresponding media sessions stored by the Phone Proxy.
- **show phone-proxy secure-phones count**—Displays the number of phones capable of secure mode stored in the database.
- **show route**—Displays the routing table.
- **show xlate count**—Shows the number of NAT sessions (xlates).

Information About the Message Severity Threshold

When you subscribe a destination profile to certain alert groups, you can set a threshold for sending alert group messages based on the message severity level. Any message with a value lower than the destination profile's specified threshold is not sent to the destination.

Table 47-1 shows the mapping between message severity levels and syslog severity levels.

Table 47-1 Message Severity Level and Syslog Level Mapping

Level	Message Severity Level	Syslog Severity Level	Description
9	Catastrophic	N/A	Network-wide catastrophic failure.
8	Disaster	N/A	Significant network impact.
7	Determined by the specified CLI keyword: subscribe-to-alert-group <i>name of alert group</i> severity <i>severity level</i>	0	Emergency. System is unusable.
6		1	Alert. Critical conditions; immediate attention needed.
5		2	Critical. Major conditions.
4		3	Error. Minor conditions.
3	Warning	4	Warning conditions.
2	Notification	5	Basic notification and informational messages. Possibly independently insignificant.
1	Normal	6	Information. Normal event, signifying a return to normal state.
0	Debugging	7	Debugging messages (default setting).

Information About Subscription Profiles

A subscription profile allows you to associate the destination recipients with interested groups. When an event registered with a subscribed group in a profile is triggered, the message associated with the event is sent to the configured recipients. Subscription profiles have the following attributes:

- You can create and configure multiple profiles.
- A profile may configure multiple e-mail or HTTPS recipients.
- A profile may subscribe multiple groups to a specified severity level.

- A profile supports three message formats: short text, long text, and XML.
- You can enable and disable a specific profile. Profiles are disabled by default.
- You can specify the maximum message size. The default is 3 MB.

A default profile, “Cisco TAC,” has been provided. The default profile has a predefined set of groups (diagnostic, environment, inventory, configuration, and telemetry) to monitor and predefined destination e-mail and HTTPS URLs. The default profile is created automatically when you initially configure Smart Call Home. The destination e-mail is `callhome@cisco.com` and the destination URL is `https://tools.cisco.com/its/service/oddce/services/DDCEService`.

**Note**

You cannot change the destination e-mail or the destination URL of the default profile.

When you subscribe a destination profile to the configuration, inventory, telemetry, or snapshot alert groups, you can choose to receive the alert group messages asynchronously or periodically at a specified time.

[Table 47-2](#) maps the default alert group to its severity level subscription and period (if applicable):

Table 47-2 Alert Group to Severity Level Subscription Mapping

Group	Severity Level	Period
Configuration	Informational	Monthly
Diagnostic	Informational and higher	N/A
Environment	Notification and higher	N/A
Inventory	Informational	Monthly
Snapshot	Informational	N/A
Syslog	Equivalent syslog	N/A
Telemetry	Informational	Daily
Test	N/A	N/A
Threat	Notification	N/A

Configuring the Environment and Snapshot Alert Groups

To configure the environment and snapshot alert groups, enter the following command:

Command	Purpose
<code>alert-group-config {environment snapshot}</code>	Enters alert-group-configuration mode.
Example: <code>hostname(config)# alert-group-config environment</code>	

Configuring Alert Group Subscription

To subscribe a destination profile to an alert group, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<code>call-home</code> Example: <code>hostname(config)# call-home</code>	Enters call-home configuration mode.
Step 2	<code>alert-group {all configuration diagnostic environment inventory syslog}</code> Example: <code>ciscoasa(cfg-call-home)# alert-group syslog</code>	Enables the specified Smart Call Home alert group. Use the all keyword to enable all alert groups. By default, all alert groups are enabled.
Step 3	<code>profile profile-name</code> Example: <code>ciscoasa(cfg-call-home)# profile CiscoTAC-1</code>	Enters the profile configuration submode for the specified destination profile. Note This is the same profile that you used in the “Enabling Smart Call Home” section on page 47-6 .
Step 4	<code>subscribe-to-alert-group all</code> Example: <code>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group all</code>	Subscribes to all available alert groups.
Step 5	<code>subscribe-to-alert-group configuration periodic {daily hh:mm monthly date hh:mm weekly day hh:mm}</code> Example: <code>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group configuration periodic weekly Wednesday 23:30</code>	Subscribes this destination profile to the configuration alert group. The periodic keyword configures the configuration alert group for periodic notification. The default period is daily. The daily keyword specifies the time of the day to send, in the <i>hh:mm</i> format, with a 24-hour clock (for example, 14:30). The weekly keyword specifies the day of the week and time of day in the <i>day hh:mm</i> format, where the day of the week is spelled out (for example, Monday). The monthly keyword specifies the numeric date, from 1 to 31, and the time of day, in the <i>date hh:mm</i> format.

	Command	Purpose
Step 6	<pre>subscribe-to-alert-group environment [severity] {catastrophic disaster emergencies alert critical errors warnings notifications informational debugging}</pre> <p>Example:</p> <pre>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group environment severity critical</pre>	<p>Subscribes to environment events with the specified optional severity level.</p> <p>The severity keyword filters messages based on the severity level, as described in Table 47-1. The default severity level is 6 (informational).</p>
Step 7	<pre>subscribe-to-alert-group syslog [severity] {catastrophic disaster fatal critical major minor warning notification normal debugging} [pattern string]</pre> <p>Example:</p> <pre>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group syslog severity notification pattern UPDOWN</pre>	<p>Subscribes to syslog events with an optional severity level or message ID.</p> <p>The severity keyword filters messages based on the severity level, as described in Table 47-1. The default severity level is 6 (informational).</p> <p>The pattern string keyword argument pair is available only if you specify the optional syslog severity level or message ID.</p>
Step 8	<pre>subscribe-to-alert-group inventory periodic {daily hh:mm monthly date hh:mm weekly day hh:mm}</pre> <p>Example:</p> <pre>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group inventory periodic daily 06:30</pre>	<p>Subscribes to inventory periodic events. The default period is daily.</p> <p>The daily keyword specifies the time of the day to send, in the <i>hh:mm</i> format, with a 24-hour clock (for example, 14:30).</p> <p>The weekly keyword specifies the day of the week and time of day in the <i>day hh:mm</i> format, where the day of the week is spelled out (for example, Monday).</p> <p>The monthly keyword specifies the numeric date, from 1 to 31, and the time of day, in the <i>date hh:mm</i> format.</p>

	Command	Purpose
Step 9	<pre>subscribe-to-alert-group telemetry periodic {hourly daily monthly day weekly day [hh:mm]}</pre> <p>Example:</p> <pre>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group monthly 15</pre>	<p>Subscribes to telemetry periodic events. The default period is daily.</p> <p>The daily keyword specifies the time of the day to send, in the <i>hh:mm</i> format, with a 24-hour clock (for example, 14:30).</p> <p>The weekly keyword specifies the day of the week and time of day in the <i>day hh:mm</i> format, where the day of the week is spelled out (for example, Monday).</p> <p>The monthly keyword specifies the numeric date, from 1 to 31, and the time of day, in the <i>date hh:mm</i> format.</p>
Step 10	<pre>subscribe-to-alert-group snapshot periodic {interval minutes hourly daily monthly day_of_month weekly day_of_week [hh:mm]}</pre> <p>Example:</p> <pre>ciscoasa(cfg-call-home-profile)# subscribe-to-alert-group snapshot periodic interval weekly wednesday 23:15</pre>	<p>Subscribes to snapshot periodic events. The default period is daily.</p> <p>The interval keyword specifies the notification interval.</p> <p>The daily keyword specifies the time of the day to send, in the <i>hh:mm</i> format, with a 24-hour clock (for example, 14:30).</p> <p>The weekly keyword specifies the day of the week and time of day in the <i>day hh:mm</i> format, where the day of the week is spelled out (for example, Monday).</p> <p>The monthly keyword specifies the numeric date, from 1 to 31, and the time of day, in the <i>date hh:mm</i> format.</p>

Optional Configuration Procedures

This section includes the following topics:

- [Configuring Smart Call Home Customer Contact Information, page 47-15](#)
- [Configuring the Mail Server, page 47-17](#)
- [Configuring Call Home Traffic Rate Limiting, page 47-18](#)
- [Testing Smart Call Home Communications, page 47-18](#)
- [Managing a Destination Profile, page 47-19](#)

Configuring Smart Call Home Customer Contact Information

You have already configured the customer e-mail address as part of the “[Enabling Smart Call Home](#)” section on [page 47-6](#). This section describes how to configure additional optional customer contact information. You can specify one or more of the following:

- Phone number
- Street address
- Customer Contract ID

- Customer name
- Cisco Customer ID
- Customer Site ID

To configure customer contact information, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home	Enters call-home configuration mode.
	Example: hostname(config)# call-home	
Step 2	(Optional) phone-number <i>phone-number-string</i>	Specifies the customer phone number. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# phone-number 8005551122	
Step 3	(Optional) street-address <i>street-address</i>	Specifies the customer address, which is a free-format string that can be up to 255 characters long. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# street-address "1234 Any Street, Any city, Any state, 12345"	
Step 4	(Optional) contact-name <i>contact-name</i>	Specifies the customer name, which can be up to 128 characters long. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# contact-name contactname1234	
Step 5	(Optional) customer-id <i>customer-id-string</i>	Specifies the Cisco customer ID, which can be up to 64 characters long. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# customer-id customer1234	
Step 6	(Optional) site-id <i>site-id-string</i>	Specifies the customer site ID, which can be up to 64 characters long. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# site-id site1234	
Step 7	(Optional) contract-id <i>contract-id-string</i>	Specifies the customer contract identification, which can be up to 128 characters long. Spaces are allowed, but you must use quotes around the string if it includes spaces.
	Example: ciscoasa(cfg-call-home)# contract-id contract1234	

Example

The following example shows how to configure contact information:

```
hostname(config)# call-home
ciscoasa(cfg-call-home)# contact-email-addr username@example.com
ciscoasa(cfg-call-home)# phone-number 8005551122
```

```

ciscoasa(cfg-call-home)# street-address "1234 Any Street, Any city, Any state, 12345"
ciscoasa(cfg-call-home)# contact-name contactname1234
ciscoasa(cfg-call-home)# customer-id customer1234
ciscoasa(cfg-call-home)# site-id site1234
ciscoasa(cfg-call-home)# contract-id contract1234

```

Configuring the Mail Server

We recommend that you use HTTPS for message transport because it is the most secure. However, you can configure an e-mail destination for Smart Call Home and then configure the mail server to use the e-mail message transport.

To configure the mail server, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home Example: hostname(config)# call-home	Enters call-home configuration mode.
Step 2	mail-server <i>ip-address</i> name priority [1-100] [all] Example: ciscoasa(cfg-call-home)# mail-server 10.10.1.1 smtp.example.com priority 1	Specifies the SMTP mail server. You can specify up to five mail servers, using five separate commands. You must configure at least one mail server for using e-mail transport of Smart Call Home messages. The lower the number, the higher the priority of the mail server. The <i>ip-address</i> argument can be an IPv4 or IPv6 mail server address.

Example

The following example shows how to configure a primary mail server (named "smtp.example.com") and a secondary mail server at IP address 10.10.1.1:

```

hostname(config)# call-home
ciscoasa(cfg-call-home)# mail-server smtp.example.com priority 1
ciscoasa(cfg-call-home)# mail-server 10.10.1.1 priority 2
ciscoasa(cfg-call-home)# exit
hostname(config)#

```

Configuring Call Home Traffic Rate Limiting

You can configure this optional setting to specify the number of messages that Smart Call Home sends per minute.

To configure Smart Call Home traffic rate limiting, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home	Enters call-home configuration mode.
	Example: hostname(config)# call-home	
Step 2	rate-limit <i>msg-count</i>	Specifies the number of messages that Smart Call Home can send per minute. The default value is 10 messages per minute.
	Example: ciscoasa(cfg-call-home)# rate-limit 5	

Example

The following example shows how to configure Smart Call Home traffic rate limiting:

```
hostname(config)# call-home
ciscoasa(cfg-call-home)# rate-limit 5
```

Testing Smart Call Home Communications

You can optionally test Smart Call Home communications by sending messages manually using two command types.

To manually send a Smart Call Home test message, enter the following command:

Command	Purpose
call-home test [<i>test-message</i>] profile <i>profile-name</i>	Sends a test message using a profile configuration.
Example: ciscoasa# call-home test [testing123] profile CiscoTAC-1	

To manually trigger a Call Home alert group message, enter the following command:

Command	Purpose
<pre>call-home send alert-group {inventory configuration snapshot telemetry} [profile profile-name]</pre> <p>Example: ciscoasa# call-home send alert-group inventory</p>	<p>Sends an alert group message to one destination profile, if specified. If no profile is specified, sends messages to all profiles that are subscribed to the inventory, configuration, snapshot, or telemetry alert groups.</p>

To execute a CLI command and e-mail the command output to Cisco TAC or to an e-mail address that you specify, enter the following command:

Command	Purpose
<pre>call-home send cli command [email email]</pre> <p>Example: ciscoasa# call-home send cli destination email username@example.com</p>	<p>Sends command output to an e-mail address. The specified CLI command can be any command, including commands for all registered modules.</p> <p>If you specify an e-mail address, the command output is sent to that address. If no e-mail address is specified, the output is sent to Cisco TAC. The e-mail is sent in log text format with the service number, if specified, in the subject line.</p> <p>The service number is required only if no e-mail address is specified, or if a Cisco TAC e-mail address is specified.</p>

Managing a Destination Profile

This section includes the following topics:

- [Configuring a Destination Profile, page 47-20](#)
- [Copying a Destination Profile, page 47-21](#)
- [Renaming a Destination Profile, page 47-21](#)

Configuring a Destination Profile

To configure a destination profile for e-mail or for HTTP, perform the following steps:

Detailed Steps

<p>Step 1</p> <p><code>call-home</code></p> <p>Example: <pre>hostname(config)# call-home</pre></p>	<p>Enters call-home configuration mode.</p>
<p>Step 2</p> <p><code>profile profile-name</code></p> <p>Example: <pre>ciscoasa(cfg-call-home)# profile newprofile</pre></p>	<p>Enters the profile configuration mode for the specified destination profile. If the specified destination profile does not exist, it is created.</p> <p>You can create a maximum of 10 active profiles. The default profile is to report back to Cisco TAC. If you want to send call home information to a different location (for example, your own server), you can configure a separate profile.</p>
<p>Step 3</p> <p><code>destination {email address http url} message-size-limit size preferred-msg-format {long-text short-text xml} transport-method {email http}</code></p> <p>Example: <pre>ciscoasa(cfg-call-home-profile)# destination address email username@example.com</pre></p> <pre>ciscoasa(cfg-call-home-profile)# destination preferred-msg-format long-text</pre>	<p>Configures the destination, message size, message format, and transport method for the smart call-home message receiver. The default message format is XML, and the default enabled transport method is e-mail. The e-mail-address is the e-mail address of the smart call-home message receiver, which can be up to 100 characters long. By default, the maximum URL size is 5 MB.</p> <p>Use the short-text format to send and read a message on a mobile device, and use the long text format to send and read a message on a computer.</p> <p>If the message receiver is the Smart Call Home back-end server, ensure that the preferred-msg-format value is XML because the back-end server can accept messages in XML format only.</p> <p>The “Enabling Smart Call Home” section on page 47-6 specifies how to set the transport method to HTTP. You can use this command to change the transport method back to e-mail.</p>

Copying a Destination Profile

To create a new destination profile by copying an existing profile, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home Example: ciscoasa(config)# call-home	Enters call-home configuration mode.
Step 2	profile profile-name Example: ciscoasa(cfg-call-home)# profile newprofile	Specifies the profile to copy.
Step 3	copy profile src-profile-name dest-profile-name Example: ciscoasa(cfg-call-home)# copy profile newprofile profile1	Copies the content of an existing profile (<i>src-profile-name</i> , which can be up to 23 characters long) to a new profile (<i>dest-profile-name</i> , which can be up to 23 characters long).

Example

The following example shows how to copy an existing profile:

```
hostname(config)# call-home
ciscoasa(cfg-call-home)# profile newprofile
ciscoasa(cfg-call-home-profile)# copy profile newprofile profile1
```

Renaming a Destination Profile

To change the name of an existing profile, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	call-home Example: hostname(config)# call-home	Enters call-home configuration mode.

	Command	Purpose
Step 2	<code>profile profilename</code> Example: ciscoasa(cfg-call-home)# profile newprofile	Specifies the profile to rename.
Step 3	<code>rename profile src-profile-name dest-profile-name</code> Example: ciscoasa(cfg-call-home)# rename profile newprofile profile1	Changes the name of an existing profile, the <i>src-profile-name</i> (an existing profile name can be up to 23 characters long), and the <i>dest-profile-name</i> (a new profile name can be up to 23 characters long).

Example

The following example shows how to rename an existing profile:

```
hostname(config)# call-home
ciscoasa(cfg-call-home)# profile newprofile
ciscoasa(cfg-call-home-profile)# rename profile newprofile profile1
```

Monitoring Anonymous Reporting and Smart Call Home

To monitor the Anonymous Reporting and Smart Call Home features, enter one of the following commands:

Command	Purpose
<code>show call-home detail</code>	Shows the current Smart Call Home detail configuration.
<code>show call-home mail-server status</code>	Shows the current mail server status.
<code>show call-home profile {profile name all}</code>	Shows the configuration of Smart Call Home profiles.
<code>show call-home registered-module status [all]</code>	Shows the registered module status.
<code>show call-home statistics</code>	Shows call-home detail status.
<code>show call-home</code>	Shows the current Smart Call Home configuration.
<code>show running-config call-home</code>	Shows the current Smart Call Home running configuration.

Command	Purpose
<code>show smart-call-home alert-group</code>	Shows the current status of Smart Call Home alert groups.
<code>show running-config all</code>	Shows details about the Anonymous Reporting user profile.

Configuration Example for Smart Call Home

The following example shows how to configure the Smart Call Home feature:

```
ciscoasa (config)# service call-home
ciscoasa (config)# call-home
ciscoasa (cfg-call-home)# contact-email-addr customer@example.com
ciscoasa (cfg-call-home)# profile CiscoTAC-1
ciscoasa (cfg-call-home-profile)# destination address http
https://example.cisco.com/its/service/example/services/ExampleService
ciscoasa (cfg-call-home-profile)# destination address email callhome@example.com
ciscoasa (cfg-call-home-profile)# destination transport-method http
ciscoasa (cfg-call-home-profile)# subscribe-to-alert-group inventory periodic daily 23:30
ciscoasa (cfg-call-home-profile)# subscribe-to-alert-group configuration periodic weekly
Wednesday 23:30
ciscoasa (cfg-call-home-profile)# subscribe-to-alert-group environment
ciscoasa (cfg-call-home-profile)# subscribe-to-alert-group diagnostic
ciscoasa (cfg-call-home-profile)# subscribe-to-alert-group telemetry periodic weekly
Monday 23:30
```

Feature History for Anonymous Reporting and Smart Call Home

Table 47-3 lists each feature change and the platform release in which it was implemented.

Table 47-3 Feature History for Anonymous Reporting and Smart Call Home

Feature Name	Platform Releases	Feature Information
Smart Call Home	8.2(2)	<p>The Smart Call Home feature offers proactive diagnostics and real-time alerts on the ASA, and provides higher network availability and increased operational efficiency.</p> <p>We introduced or modified the following commands:</p> <p>active (call home), call-home, call-home send alert-group, call-home test, contact-email-addr, customer-id (call home), destination (call home), profile, rename profile, service call-home, show call-home, show call-home detail, show smart-call-home alert-group, show call-home profile, show call-home statistics, show call-home mail-server status, show running-config call-home, show call-home registered-module status all, site-id, street-address, subscribe-to-alert-group all, alert-group-config, subscribe-to-alert-group configuration, subscribe-to-alert-group diagnostic, subscribe-to-alert-group environment, subscribe-to-alert-group inventory periodic, subscribe-to-alert-group snapshot periodic, subscribe-to-alert-group syslog, subscribe-to-alert-group telemetry periodic.</p>
Anonymous Reporting	9.0(1)	<p>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.</p> <p>We introduced the following commands: call-home reporting anonymous, call-home test reporting anonymous.</p>

Table 47-3 Feature History for Anonymous Reporting and Smart Call Home (continued)

Feature Name	Platform Releases	Feature Information
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)	<p>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:</p> <ul style="list-style-type: none"> • When a unit joins the cluster • When a unit leaves the cluster • When a cluster unit becomes the cluster master <p>Each message that is sent includes the following information:</p> <ul style="list-style-type: none"> • The active cluster member count • The output of the show cluster info command and the show cluster history command on the cluster master



PART 10

Reference



Using the Command-Line Interface

This appendix describes how to use the CLI on the ASA and includes the following sections:

- [Firewall Mode and Security Context Mode, page 48-1](#)
- [Command Modes and Prompts, page 48-2](#)
- [Syntax Formatting, page 48-3](#)
- [Abbreviating Commands, page 48-3](#)
- [Command-Line Editing, page 48-3](#)
- [Command Completion, page 48-4](#)
- [Command Help, page 48-4](#)
- [Viewing the Running Configuration, page 48-4](#)
- [Filtering show and more Command Output, page 48-5](#)
- [Command Output Paging, page 48-5](#)
- [Adding Comments, page 48-6](#)
- [Text Configuration Files, page 48-6](#)
- [Supported Character Sets, page 48-8](#)



Note

The CLI uses similar syntax and other conventions to the Cisco IOS CLI, but the ASA operating system is not a version of Cisco IOS software. Do not assume that a Cisco IOS CLI command works with or has the same function on the ASA.

Firewall Mode and Security Context Mode

The ASA runs in a combination of the following modes:

- Transparent firewall or routed firewall mode

The firewall mode determines if the ASA runs as a Layer 2 or Layer 3 firewall.

- Multiple context or single context mode

The security context mode determines if the ASA runs as a single device or as multiple security contexts, which act like virtual devices.

Some commands are only available in certain modes.

Command Modes and Prompts

The ASA CLI includes command modes. Some commands can only be entered in certain modes. For example, to enter commands that show sensitive information, you need to enter a password and enter a more privileged mode. Then, to ensure that configuration changes are not entered accidentally, you have to enter a configuration mode. All lower commands can be entered in higher modes, for example, you can enter a privileged EXEC command in global configuration mode.



Note

The various types of prompts are all default prompts and when configured, they can be different.

- When you are in the system configuration or in single context mode, the prompt begins with the hostname:

```
ciscoasa
```
- When printing the prompt string, the prompt configuration is parsed and the configured keyword values are printed in the order in which you have set the **prompt** command. The keyword arguments can be any of the following and in any order: hostname, domain, context, priority, state.

```
asa(config)# prompt hostname context priority state
```
- When you are within a context, the prompt begins with the hostname followed by the context name:

```
ciscoasa/context
```

The prompt changes depending on the access mode:

- User EXEC mode
 User EXEC mode lets you see minimum ASA settings. The user EXEC mode prompt appears as follows when you first access the ASA:

```
ciscoasa>
```

```
ciscoasa/context>
```
- Privileged EXEC mode
 Privileged EXEC mode lets you see all current settings up to your privilege level. Any user EXEC mode command will work in privileged EXEC mode. Enter the **enable** command in user EXEC mode, which requires a password, to start privileged EXEC mode. The prompt includes the number sign (#):

```
ciscoasa#
```

```
ciscoasa/context#
```
- Global configuration mode
 Global configuration mode lets you change the ASA configuration. All user EXEC, privileged EXEC, and global configuration commands are available in this mode. Enter the **configure terminal** command in privileged EXEC mode to start global configuration mode. The prompt changes to the following:

```
ciscoasa(config)#
```

```
ciscoasa/context(config)#
```
- Command-specific configuration modes

From global configuration mode, some commands enter a command-specific configuration mode. All user EXEC, privileged EXEC, global configuration, and command-specific configuration commands are available in this mode. For example, the **interface** command enters interface configuration mode. The prompt changes to the following:

```
ciscoasa(config-if)#
ciscoasa/context(config-if)#
```

Syntax Formatting

Command syntax descriptions use the conventions listed in [Table 48-1](#).

Table 48-1 *Syntax Conventions*

Convention	Description
bold	Bold text indicates commands and keywords that you enter literally as shown.
<i>italics</i>	Italic text indicates arguments for which you supply values.
[x]	Square brackets enclose an optional element (keyword or argument).
	A vertical bar indicates a choice within an optional or required set of keywords or arguments.
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested sets of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.

Abbreviating Commands

You can abbreviate most commands down to the fewest unique characters for a command; for example, you can enter **wr t** to view the configuration instead of entering the full command **write terminal**, or you can enter **en** to start privileged mode and **conf t** to start configuration mode. In addition, you can enter **0** to represent **0.0.0.0**.

Command-Line Editing

The ASA uses the same command-line editing conventions as Cisco IOS software. You can view all previously entered commands with the **show history** command or individually with the up arrow or **^p** command. Once you have examined a previously entered command, you can move forward in the list with the down arrow or **^n** command. When you reach a command you wish to reuse, you can edit it or press the **Enter** key to start it. You can also delete the word to the left of the cursor with **^w**, or erase the line with **^u**.

The ASA permits up to 512 characters in a command; additional characters are ignored.

Command Completion

To complete a command or keyword after entering a partial string, press the **Tab** key. The ASA only completes the command or keyword if the partial string matches only one command or keyword. For example, if you enter `s` and press the **Tab** key, the ASA does not complete the command because it matches more than one command. However, if you enter `dis`, the **Tab** key completes the `disable` command.

Command Help

Help information is available from the command line by entering the following commands:

- `help command_name`
Shows help for the specific command.
- `command_name ?`
Shows a list of arguments available.
- `string?` (no space)
Lists the possible commands that start with the string.
- `? and +?`
Lists all commands available. If you enter `?`, the ASA shows only commands available for the current mode. To show all commands available, including those for lower modes, enter `+?`.



Note

If you want to include a question mark (?) in a command string, you must press **Ctrl-V** before typing the question mark so that you do not inadvertently invoke CLI help.

Viewing the Running Configuration

To view the running configuration, use one of the following commands.

To filter the command output, see the [“Filtering show and more Command Output”](#) section on page 48-5.

Command	Purpose
<code>show running-config [all] [command]</code>	Shows the running configuration. If you specify all , then all default settings are shown as well. If you specify a <i>command</i> , then the output only includes related commands. Note Many passwords are shown as <code>*****</code> . To view the passwords in plain text, or in encrypted form if you have a master passphrase enabled (see the “Configuring the Master Passphrase” section on page 13-8), use the more command below.
<code>more system:running-config</code>	Shows the running configuration. Passwords are shown in plain text or in encrypted form if you have a master passphrase enabled (see the “Configuring the Master Passphrase” section on page 13-8).

Filtering show and more Command Output

You can use the vertical bar (|) with any **show** command and include a filter option and filtering expression. The filtering is performed by matching each output line with a regular expression, similar to Cisco IOS software. By selecting different filter options you can include or exclude all output that matches the expression. You can also display all output beginning with the line that matches the expression.

The syntax for using filtering options with the **show** command is as follows:

```
ciscoasa# show command | {include | exclude | begin | grep [-v]} regex
```

or

```
ciscoasa# more system:running-config | {include | exclude | begin | grep [-v]} regex
```



Note

The **more** command can view the contents of any file, not just the running configuration; see the command reference for more information.

In this command string, the first vertical bar (|) is the operator and must be included in the command. This operator directs the output of the **show** command to the filter. In the syntax diagram, the other vertical bars (|) indicate alternative options and are not part of the command.

The **include** option includes all output lines that match the regular expression. The **grep** option without **-v** has the same effect. The **exclude** option excludes all output lines that match the regular expression. The **grep** option with **-v** has the same effect. The **begin** option shows all the output lines starting with the line that matches the regular expression.

Replace *regex* with any Cisco IOS regular expression. The regular expression is not enclosed in quotes or double-quotes, so be careful with trailing white spaces, which will be taken as part of the regular expression.

When creating regular expressions, you can use any letter or number that you want to match. In addition, certain keyboard characters called *metacharacters* have special meaning when used in regular expressions.

Use **Ctrl+V** to escape all of the special characters in the CLI, such as a question mark (?) or a tab. For example, type **d[Ctrl+V]?g** to enter **d?g** in the configuration.

For a list of metacharacters, see [Table 17-1 on page 17-15](#).

Command Output Paging

For commands such as **help** or **?**, **show**, **show xlate**, or other commands that provide long listings, you can determine if the information displays a screen and pauses, or lets the command run to completion. The **pager** command lets you choose the number of lines to display before the More prompt appears.

When paging is enabled, the following prompt appears:

```
<--- More --->
```

The More prompt uses syntax similar to the UNIX **more** command:

- To view another screen, press the **Space** bar.
- To view the next line, press the **Enter** key.
- To return to the command line, press the **q** key.

Adding Comments

You can precede a line with a colon (:) to create a comment. However, the comment only appears in the command history buffer and not in the configuration. Therefore, you can view the comment with the **show history** command or by pressing an arrow key to retrieve a previous command, but because the comment is not in the configuration, the **write terminal** command does not display it.

Text Configuration Files

This section describes how to format a text configuration file that you can download to the ASA, and includes the following topics:

- [How Commands Correspond with Lines in the Text File, page 48-6](#)
- [Command-Specific Configuration Mode Commands, page 48-6](#)
- [Automatic Text Entries, page 48-7](#)
- [Line Order, page 48-7](#)
- [Commands Not Included in the Text Configuration, page 48-7](#)
- [Passwords, page 48-7](#)
- [Multiple Security Context Files, page 48-7](#)

How Commands Correspond with Lines in the Text File

The text configuration file includes lines that correspond with the commands described in this guide.

In examples, commands are preceded by a CLI prompt. The prompt in the following example is “ciscoasa(config)#”:

```
ciscoasa(config)# context a
```

In the text configuration file you are not prompted to enter commands, so the prompt is omitted:

```
context a
```

Command-Specific Configuration Mode Commands

Command-specific configuration mode commands appear indented under the main command when entered at the command line. Your text file lines do not need to be indented, as long as the commands appear directly following the main command. For example, the following unindented text is read the same as indented text:

```
interface gigabitethernet0/0  
nameif inside  
interface gigabitethernet0/1  
    nameif outside
```


Automatic Text Entries

When you download a configuration to the ASA, it inserts some lines automatically. For example, the ASA inserts lines for default settings or for the time the configuration was modified. You do not need to enter these automatic entries when you create your text file.

Line Order

For the most part, commands can be in any order in the file. However, some lines, such as ACEs, are processed in the order they appear, and the order can affect the function of the access list. Other commands might also have order requirements. For example, you must enter the **nameif** command for an interface first because many subsequent commands use the name of the interface. Also, commands in a command-specific configuration mode must directly follow the main command.

Commands Not Included in the Text Configuration

Some commands do not insert lines in the configuration. For example, a runtime command such as **show running-config** does not have a corresponding line in the text file.

Passwords

The login, enable, and user passwords are automatically encrypted before they are stored in the configuration. For example, the encrypted form of the password “cisco” might look like jMorNbK0514fadBh. You can copy the configuration passwords to another ASA in its encrypted form, but you cannot unencrypt the passwords yourself.

If you enter an unencrypted password in a text file, the ASA does not automatically encrypt it when you copy the configuration to the ASA. The ASA only encrypts it when you save the running configuration from the command line using the **copy running-config startup-config** or **write memory** command.

Multiple Security Context Files

For multiple security contexts, the entire configuration consists of the following multiple parts:

- The security context configurations
- The system configuration, which identifies basic settings for the ASA, including a list of contexts
- The admin context, which provides network interfaces for the system configuration

The system configuration does not include any 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 a context that is designated as the admin context.

Each context is similar to a single context mode configuration. The system configuration differs from a context configuration in that the system configuration includes system-only commands (such as a list of all contexts) while other typical commands are not present (such as many interface parameters).

Supported Character Sets

The ASA CLI currently supports UTF-8 encoding only. UTF-8 is the particular encoding scheme for Unicode symbols, and has been designed to be compatible with an ASCII subset of symbols. ASCII characters are represented in UTF-8 as one-byte characters. All other characters are represented in UTF-8 as multibyte symbols.

The ASCII printable characters (0x20 to 0x7e) are fully supported. The printable ASCII characters are the same as ISO 8859-1. UTF-8 is a superset of ISO 8859-1, so the first 256 characters (0-255) are the same as ISO 8859-1. The ASA CLI supports up to 255 characters (multibyte characters) of ISO 8859-1.



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 49-1](#)
- [IPv6 Addresses, page 49-5](#)
- [Protocols and Applications, page 49-11](#)
- [TCP and UDP Ports, page 49-11](#)
- [Local Ports and Protocols, page 49-14](#)
- [ICMP Types, page 49-15](#)

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 49-1](#)
- [Private Networks, page 49-2](#)
- [Subnet Masks, page 49-2](#)

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 49-3](#)
- [Determining the Address to Use with the Subnet Mask, page 49-3](#)

Determining the Subnet Mask

To determine the subnet mask based on how many hosts you want, see [Table 49-1](#).

Table 49-1 *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 49-3](#)
- [Class B-Size Network Address, page 49-4](#)

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 49-2](#) shows the 8-host subnets (/29) of 192.168.0.x.

Table 49-2 *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 49-2 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 49-3](#) shows the 16 subnets of the network 10.1.

Table 49-3 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 49-5](#)
- [IPv6 Address Types, page 49-6](#)
- [IPv6 Address Prefixes, page 49-10](#)



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 11-12](#)

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. 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 49-4](#) shows several examples of address compression for different types of IPv6 address.

Table 49-4 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	FF01:0:0:0:0:0:101	FF01::101
Loopback	0:0:0:0:0:0:1	::1
Unspecified	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: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: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 49-6](#)
- [Multicast Address, page 49-8](#)
- [Anycast Address, page 49-9](#)
- [Required Addresses, page 49-10](#)

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 49-7](#)
- [Site-Local Address, page 49-7](#)
- [Link-Local Address, page 49-7](#)
- [IPv4-Compatible IPv6 Addresses, page 49-7](#)
- [Unspecified Address, page 49-8](#)
- [Loopback Address, page 49-8](#)
- [Interface Identifiers, page 49-8](#)

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 49-10, 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 49-8, 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 49-7).

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, 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, 0:0:0:0:0:0:1, 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.

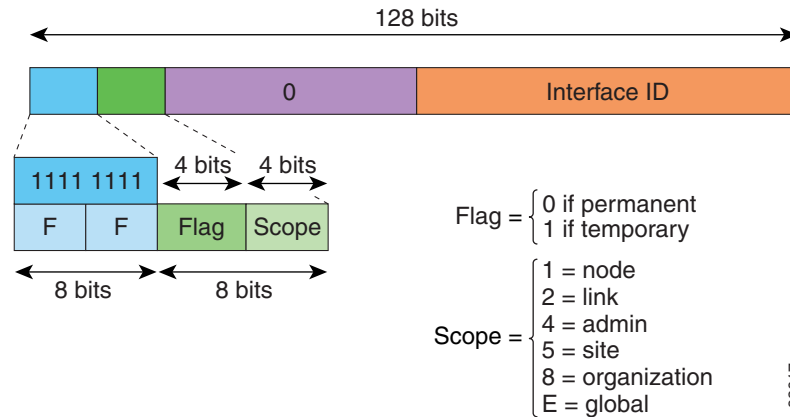
For example, an interface with the MAC address of 00E0.b601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

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 49-1](#) shows the format of the IPv6 multicast address.

Figure 49-1 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: FF02:0:0:0:1: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 49-5](#) shows the prefixes for each IPv6 address type.

Table 49-5 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 49-6 lists the protocol literal values and port numbers; either can be entered in ASA commands.

Table 49-6 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 49-7 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 49-7 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 49-7 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 49-7 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 49-8 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 49-8, 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 49-8 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 49-8 *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 49-9 lists the ICMP type numbers and names that you can enter in ASA commands.

Table 49-9 *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 49-9 ICMP Types (continued)

ICMP Number	ICMP Name
31	conversion-error
32	mobile-redirect



Symbols

- ?
 - command string [48-4](#)
 - help [48-4](#)
- /bits subnet masks [49-3](#)

Numerics

- 4GE SSM
 - connector types [9-15](#)
 - fiber [9-15](#)
 - SFP [9-15](#)
- 802.1Q tagging [10-9](#)
- 802.1Q trunk [9-33](#)

A

AAA

- about [32-1, 33-1, 34-1, 35-1, 37-1](#)
- authentication
 - CLI access [41-20](#)
 - privileged EXEC mode [41-21](#)
- authorization
 - command [41-27](#)
- server [44-4](#)
 - adding [34-15, 36-7, 37-3, 37-4](#)
 - types [32-1](#)
- support summary [32-3](#)
- abbreviating commands [48-3](#)
- ABR
 - definition of [27-2](#)
- Access Group pane

- description [30-8](#)
- access lists
 - about [18-1](#)
 - ACE logging, configuring [23-1](#)
 - deny flows, managing [23-5](#)
 - implicit deny [18-3](#)
 - IP address guidelines [18-3](#)
 - logging [23-1](#)
 - NAT guidelines [18-3](#)
 - remarks [19-9](#)
 - scheduling activation [19-2](#)
 - types [18-1](#)
- access ports [10-7](#)
- ACEs
 - See* access lists
- activation key
 - entering [4-36](#)
 - location [4-34](#)
 - obtaining [4-35](#)
- Active/Active failover
 - about [7-22](#)
 - actions [7-23](#)
 - configuring
 - asymmetric routing support [7-39](#)
 - failover group preemption [7-37](#)
 - duplicate MAC addresses, avoiding [7-8](#)
 - primary status [7-22](#)
 - secondary status [7-22](#)
- Active/Standby failover
 - about [7-20](#)
 - actions [7-20](#)
 - command replication [7-19](#)
 - configuration synchronization [7-18](#)

- device initialization [7-18](#)
- primary unit [7-20](#)
- secondary unit [7-20](#)
- Adaptive Security Algorithm [1-17](#)
- Add/Edit Access Group dialog box
 - description [30-8](#)
- Add/Edit Filtering Entry dialog box
 - description [27-27](#)
- Add/Edit IGMP Join Group dialog box
 - description [30-7](#)
- 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 Summary Address dialog box
 - description [27-13, 27-19](#)
- Add/Edit Virtual Link dialog box
 - description [27-28](#)
- admin context
 - about [6-2](#)
 - changing [6-26](#)
- administrative access
 - using ICMP for [41-11](#)
- administrative distance [25-3, 25-5](#)
- Advanced OSPF Interface Properties dialog box [27-18](#)
- Advanced OSPF Virtual Link Properties dialog box
 - description [27-28](#)
- AIP SSM
 - port-forwarding
 - enabling [11-7, 12-9](#)
- alternate address, ICMP message [49-15](#)
- analyzing syslog messages [44-2](#)
- application inspection
 - security level requirements [11-2, 12-2](#)
- Area/Networks tab
 - description [27-6](#)
- area border router [27-2](#)
- ARP inspection
 - about [5-6](#)
 - enabling [5-11](#)
 - static entry [5-10](#)
- ARP spoofing [5-6](#)
- ARP test, failover [7-17](#)
- ASA (Adaptive Security Algorithm) [1-17](#)
- ASA 5505
 - Base license [10-2](#)
 - MAC addresses [10-4](#)
 - maximum VLANs [10-2](#)
 - native VLAN support [10-10](#)
 - non-forwarding interface [10-7](#)
 - power over Ethernet [10-4](#)
 - protected switch ports [10-8, 10-10](#)
 - Security Plus license [10-2](#)
 - SPAN [10-4](#)
 - Spanning Tree Protocol, unsupported [10-8](#)
- ASA 5550 throughput [11-7, 12-10](#)
- ASBR
 - definition of [27-2](#)
- ASDM software
 - allowing access [41-6](#)
 - installing [42-17](#)
- ASR [7-39](#)
- ASR groups [7-39](#)
- asymmetric routing support [7-39](#)
- attributes
 - RADIUS [34-3](#)
- attribute-value pairs
 - TACACS+ [35-1](#)
- authentication
 - about [32-1](#)
 - CLI access [41-20](#)
 - privileged EXEC mode [41-21](#)
- Authentication tab
 - description [27-16](#)
- authorization
 - about [32-2](#)
 - command [41-27](#)

Auto-MDI/MDIX [9-2, 10-4](#)

autostate messaging [2-11](#)

Auto-Update, configuring [42-36](#)

B

Baltimore Technologies, CA server support [40-4](#)

bits subnet masks [49-3](#)

BPDUs

forwarding on the switch [2-11](#)

bridge

entry timeout [5-12](#)

table, *See* MAC address table

broadcast Ping test [7-17](#)

building blocks [17-1](#)

bypassing the firewall, in the switch [2-5](#)

C

CA

CRs and [40-2](#)

public key cryptography [40-2](#)

revoked certificates [40-2](#)

supported servers [40-4](#)

capturing packets [43-2](#)

CA server

Digicert [40-4](#)

Geotrust [40-4](#)

Godaddy [40-4](#)

iPlanet [40-4](#)

Netscape [40-4](#)

RSA Keon [40-4](#)

Thawte [40-4](#)

Catalyst 6500

See switch

certificate

enrollment protocol [40-13](#)

Certificate Revocation Lists

See CRLs

change query interval [30-9](#)

change query response time [30-9](#)

change query timeout value [30-9](#)

changing between contexts [6-24](#)

changing the severity level [44-19](#)

Cisco [14-7](#)

Cisco 7600

See switch

Cisco IOS CS CA

server support [40-4](#)

Cisco IP Phones

DHCP [14-7](#)

Class A, B, and C addresses [49-1](#)

classes, logging

filtering messages by [44-17](#)

message class variables [44-4](#)

types [44-4](#)

classes, resource

See resource management

class map

regular expression [17-17](#)

CLI

abbreviating commands [48-3](#)

adding comments [48-6](#)

command line editing [48-3](#)

command output paging [48-5](#)

displaying [48-5](#)

help [48-4](#)

paging [48-5](#)

syntax formatting [48-3](#)

clustering

ASDM connection certificate IP address mismatch [8-12](#)

backup owner [8-10](#)

cabling [8-33](#)

cluster control link

configuring [8-43, 8-49](#)

failure [8-9](#)

- MTU [8-45](#)
- overview [8-7](#)
- redundancy [8-8](#)
- size [8-7](#)
- configuration
 - examples [8-63](#)
 - replication [8-11](#)
- connection
 - new, ownership [8-19](#)
 - rebalancing [8-47](#)
- console replication [8-48](#)
- context mode [8-28](#)
- data path connection state replication [8-10](#)
- device-local EtherChannels, configuring on switch [8-30](#)
- executing a command cluster-wide [8-58](#)
- failover [8-28](#)
- feature history [8-78](#)
- features
 - centralized [8-21](#)
 - individual units [8-22](#)
 - NAT [8-24](#)
 - SNMP [8-26](#)
 - syslog and netflow [8-26](#)
 - unsupported [8-20](#)
 - VPN [8-26](#)
- guidelines and limitations [8-28](#)
- high availability [8-9](#)
- individual cluster interfaces, configuring [8-36](#)
- interface monitoring [8-9](#)
- IPv6 [8-28](#)
- key [8-46, 8-52](#)
- licensing [8-27](#)
- management
 - interface, configuring [8-36](#)
 - interface, overview [8-11](#)
 - network [8-11](#)
 - overview [8-11](#)
- master unit
 - changing [8-57](#)
 - election [8-3](#)
 - maximum members [8-27](#)
 - member requirements [8-3](#)
 - model support [8-28](#)
 - monitoring [8-59](#)
 - overview
 - bootstrap configuration [8-3](#)
 - cluster control link [8-7](#)
 - Equal-Cost Multi-Path Routing [8-16](#)
 - interfaces [8-4](#)
 - load balancing [8-13](#)
 - management [8-11](#)
 - master unit [8-3](#)
 - Policy-Based Routing [8-15](#)
 - spanned EtherChannel [8-13](#)
 - performance scaling factor [8-2](#)
 - prerequisites [8-27](#)
 - rebalancing new connections [8-20](#)
 - removing a member [8-55](#)
 - RSA key replication [8-12](#)
 - software requirements [8-3](#)
 - spanned EtherChannel
 - benefits [8-13](#)
 - configuring [8-38](#)
 - load balancing [8-14](#)
 - maximum throughput [8-13](#)
 - overview [8-13](#)
 - redundancy [8-14](#)
 - VSS or vPC [8-14](#)
 - spanning-tree portfast [8-28](#)
 - unit failure [8-9](#)
 - unit health monitoring [8-9](#)
 - upgrading software [8-3](#)
- command authorization
 - about [41-16](#)
 - configuring [41-27](#)
 - multiple contexts [41-17](#)
- command prompts [48-2](#)

- comments
 - configuration [48-6](#)
 - configuration
 - clearing [3-27](#)
 - comments [48-6](#)
 - factory default
 - commands [3-18](#)
 - restoring [3-19](#)
 - saving [3-25](#)
 - switch [2-1](#)
 - text file [3-28](#)
 - URL for a context [6-22](#)
 - viewing [3-27](#)
 - configuration examples
 - logging [44-21](#)
 - configuration examples for SNMP [45-28](#)
 - configuration mode
 - accessing [3-2, 3-4](#)
 - prompt [48-2](#)
 - connection limits
 - per context [6-17](#)
 - console port logging [44-12](#)
 - context mode [28-3](#)
 - context modes [25-2, 26-3, 27-3, 29-3, 30-3](#)
 - contexts
 - See* security contexts
 - conversion error, ICMP message [49-16](#)
 - Coredump [43-6](#)
 - crash dump [43-6](#)
 - creating a custom event list [44-14](#)
 - custom messages list
 - logging output destination [44-5](#)
-
- D**
- data flow
 - routed firewall [5-14](#)
 - transparent firewall [5-20](#)
 - date and time in messages [44-19](#)
 - DDNS [15-2](#)
 - debug messages [43-1](#)
 - default
 - class [6-9](#)
 - routes, defining equal cost routes [25-4](#)
 - default configuration
 - commands [3-18](#)
 - restoring [3-19](#)
 - default routes
 - about [25-4](#)
 - configuring [25-4](#)
 - delay sending flow-create events
 - flow-create events
 - delay sending [46-9](#)
 - deleting files from Flash [42-12](#)
 - deny flows, logging [23-5](#)
 - device ID, including in messages [44-18](#)
 - device ID in messages [44-18](#)
 - DHCP
 - Cisco IP Phones [14-7](#)
 - options [14-6](#)
 - relay [14-8](#)
 - server [14-5](#)
 - DHCP Relay panel [15-9](#)
 - DHCP services [13-8](#)
 - directory hierarchy search [36-3](#)
 - disabling messages [44-19](#)
 - disabling messages, specific message IDs [44-19](#)
 - DMZ, definition [1-14](#)
 - DNS
 - server, configuring [13-13](#)
 - domain name [13-4](#)
 - dotted decimal subnet masks [49-3](#)
 - dual IP stack, configuring [11-2](#)
 - dual-ISP support [25-6](#)
 - duplex, configuring [9-15, 10-6](#)

E

echo reply, ICMP message [49-15](#)

ECMP [25-3](#)

editing command lines [48-3](#)

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 [28-15](#)
- hello packets [28-1](#)
- hold time [28-2, 28-15](#)
- neighbor discovery [28-1](#)
- stub routing [28-4](#)
- stuck-in-active [28-2](#)

enable command [3-1](#)

enabling logging [44-7](#)

enabling secure logging [44-17](#)

Entrust, CA server support [40-4](#)

established command, security level requirements [11-2, 12-2](#)

EtherChannel

- adding interfaces [9-30](#)
- channel group [9-30](#)
- compatibility [9-5](#)
- converting existing interfaces [9-16](#)
- example [9-37](#)
- failover [9-13](#)
- guidelines [9-13](#)
- interface requirements [9-5](#)

LACP [9-6](#)

load balancing

- configuring [9-32](#)
- overview [9-7](#)

MAC address [9-8](#)

management interface [9-30](#)

maximum interfaces [9-32](#)

minimum interfaces [9-32](#)

mode

- active [9-7](#)
- on [9-7](#)
- passive [9-7](#)

monitoring [9-36](#)

overview [9-5](#)

port priority [9-30](#)

system priority [9-32](#)

Ethernet

- Auto-MDI/MDIX [9-2, 10-4](#)
- duplex [9-15, 10-6](#)
- jumbo frames, ASA 5580 [9-35](#)
- MTU [11-12, 12-15](#)
- speed [9-15, 10-6](#)

evaluation license [4-24](#)

exporting NetFlow records [46-5](#)

extended ACLs

- configuring
 - for management traffic [19-4](#)

F

facility, syslog [44-9](#)

factory default configuration

- commands [3-18](#)
- restoring [3-19](#)

failover

- about [7-1](#)
- Active/Active, *See* Active/Active failover
- Active/Standby, *See* Active/Standby failover
- configuration file
 - terminal messages, Active/Standby [7-18](#)
- contexts [7-20](#)
- debug messages [7-48](#)
- disabling [7-43](#)
- Ethernet failover cable [7-4](#)
- failover link [7-3](#)
- forcing [7-42](#)
- guidelines [45-17](#)

- health monitoring [7-16](#)
- interface health [7-17](#)
- interface monitoring [7-17](#)
- interface tests [7-17](#)
- link communications [7-3](#)
- MAC addresses
 - about [7-20](#)
 - automatically assigning [6-12](#)
- module placement
 - inter-chassis [7-9](#)
 - intra-chassis [7-8](#)
- monitoring, health [7-16](#)
- network tests [7-17](#)
- primary unit [7-20](#)
- redundant interfaces [9-13](#)
- restoring a failed group [7-44](#)
- restoring a failed unit [7-44](#)
- secondary unit [7-20](#)
- SNMP syslog traps [7-48](#)
- Stateful Failover, *See* Stateful Failover
- state link [7-4](#)
- switch configuration [2-11](#)
- system log messages [7-48](#)
- system requirements [7-2](#)
- testing [7-44](#)
- trunk [2-11](#)
- unit health [7-16](#)

fast path [1-18](#)

fiber interfaces [9-15](#)

Fibre Channel interfaces

- default settings [20-2, 21-2, 22-3](#)

filtering

- security level requirements [11-2, 12-2](#)
- show command output [48-5](#)

filtering messages [44-4](#)

Filtering pane

- description [27-27](#)

firewall mode

- about [5-1](#)

- configuring [5-1](#)
- Flash memory
 - removing files [42-12](#)
- flash memory available for logs [44-16](#)
- flow control for 10 Gigabit Ethernet [9-26](#)
- flow-export actions [46-4](#)
- format of messages [44-3](#)
- fragment protection [1-15](#)

G

- generating RSA keys [39-16, 39-18, 39-20, 39-22, 40-11](#)
- groups
 - SNMP [45-16](#)

H

H.323

- transparent firewall guidelines [5-6](#)

help, command line [48-4](#)

high availability

- about [7-1](#)

host

- SNMP [45-16](#)

hostname

- configuring [13-3](#)
- in banners [13-3](#)
- multiple context mode [13-3](#)

hosts, subnet masks for [49-3](#)

HSRP [5-5](#)

HTTP(S)

- authentication [41-21](#)

HTTPS/Telnet/SSH

- allowing network or host access to ASDM [41-1](#)

I

ICMP

- rules for access to ADSM [41-11](#)
- type numbers [49-15](#)
- implementing SNMP [45-16](#)
- information reply, ICMP message [49-15](#)
- information request, ICMP message [49-15](#)
- inside, definition [1-14](#)
- installation
 - module verification [2-6](#)
- interface
 - MTU [11-12, 12-15](#)
- Interface pane [27-16](#)
- interfaces
 - ASA 5505
 - enabled status [10-7](#)
 - MAC addresses [10-4](#)
 - maximum VLANs [10-2](#)
 - non-forwarding [10-7](#)
 - protected switch ports [10-8, 10-10](#)
 - switch port configuration [10-7](#)
 - trunk ports [10-9](#)
 - ASA 5550 throughput [11-7, 12-10](#)
 - default settings [20-2, 21-2, 22-3](#)
 - duplex [9-15, 10-6](#)
 - enabling [9-27](#)
 - failover monitoring [7-17](#)
 - fiber [9-15](#)
 - IDs [9-26](#)
 - IP address [11-8, 12-13](#)
 - MAC addresses
 - automatically assigning [6-24](#)
 - manually assigning to interfaces [11-11, 12-15](#)
 - mapped name [6-21](#)
 - naming, physical and subinterface [11-8, 12-11, 12-12](#)
 - redundant [9-28](#)
 - SFP [9-15](#)
 - speed [9-15, 10-6](#)
 - subinterfaces [9-33](#)
 - turning off [11-17, 12-19](#)
 - turning on [11-17, 12-19](#)

- IOS
 - upgrading [2-3](#)
- IP addresses
 - classes [49-1](#)
 - interface [11-8, 12-13](#)
 - management, transparent firewall [12-8](#)
 - private [49-2](#)
 - subnet mask [49-4](#)
- IPv6
 - configuring alongside IPv4 [11-2](#)
 - default route [25-5](#)
 - dual IP stack [11-2](#)
 - duplicate address detection [31-2](#)
 - neighbor discovery [31-1](#)
 - router advertisement messages [31-3](#)
 - static neighbors [31-4](#)
 - static routes [25-5](#)
- IPv6 addresses
 - anycast [49-9](#)
 - format [49-5](#)
 - multicast [49-8](#)
 - prefixes [49-10](#)
 - required [49-10](#)
 - types of [49-6](#)
 - unicast [49-6](#)
- IPv6 prefixes [31-12](#)
- IPX [2-5](#)

J

- Join Group pane
 - description [30-7](#)
- jumbo frames, ASA 5580 [9-35](#)

K

- Kerberos
 - configuring [34-15, 36-7, 37-3](#)

- L**
- LACP [9-6](#)
 - 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](#)
 - user authorization [36-10](#)
 - licenses
 - activation key
 - entering [4-36](#)
 - location [4-34](#)
 - obtaining [4-35](#)
 - ASA 5505 [4-3](#)
 - ASA 5510 [4-4, 4-9](#)
 - ASA 5520 [4-5](#)
 - ASA 5540 [4-6](#)
 - ASA 5550 [4-7](#)
 - ASA 5580 [4-8, 4-17](#)
 - ASA 5585-X [4-16](#)
 - default [4-24](#)
 - evaluation [4-24](#)
 - failover [4-34](#)
 - guidelines [4-33](#)
 - managing [4-1](#)
 - preinstalled [4-24](#)
 - Product Authorization Key [4-35](#)
 - shared
 - backup server, configuring [4-39](#)
 - backup server, information [4-28](#)
 - client, configuring [4-39](#)
 - communication issues [4-28](#)
 - failover [4-29](#)
 - maximum clients [4-29](#)
 - monitoring [4-49](#)
 - overview [4-27](#)
 - server, configuring [4-37](#)
 - SSL messages [4-28](#)
 - temporary [4-24](#)
 - viewing current [4-40](#)
 - VPN Flex [4-24](#)
 - licensing requirements
 - logging [44-5](#)
 - licensing requirements for SNMP [45-17](#)
 - link up/down test [7-17](#)
 - local user database
 - adding a user [33-4](#)
 - configuring [33-4](#)
 - logging in [41-22](#)
 - lockout recovery [41-36](#)
 - logging
 - access lists [23-1](#)
 - classes
 - filtering messages by [44-4](#)
 - types [44-4, 44-17](#)
 - device-id, including in system log messages [44-18](#)
 - e-mail
 - source address [44-11](#)
 - EMBLEM format [44-15](#)
 - facility option [44-9](#)
 - filtering
 - by message class [44-17](#)
 - by message list [44-5](#)
 - by severity level [44-1](#)
 - logging queue, configuring [44-16](#)
 - output destinations [44-8](#)
 - console port [44-8, 44-11, 44-12](#)
 - internal buffer [44-1, 44-7](#)
 - Telnet or SSH session [44-7](#)
 - queue
 - changing the size of [44-16](#)
 - configuring [44-16](#)

- viewing queue statistics [44-20](#)
- severity level, changing [44-20](#)
- timestamp, including [44-19](#)
- logging feature history [44-21](#)
- logging queue
 - configuring [44-16](#)
- login
 - banner, configuring [41-7](#)
 - console [3-1](#)
 - enable [3-1](#)
 - global configuration mode [3-2](#)
 - local user [41-22](#)
 - password [13-2](#)
 - session [3-4](#)
 - SSH [3-4, 41-5](#)
 - Telnet [3-4, 13-2](#)
- loops, avoiding [2-11](#)

M

- MAC address
 - redundant interfaces [9-5](#)
- MAC addresses
 - ASA 5505 [10-4](#)
 - automatically assigning [6-24](#)
 - failover [7-20](#)
 - manually assigning to interfaces [11-11, 12-15](#)
 - security context classification [6-3](#)
- MAC address table
 - about [5-20](#)
 - built-in-switch [5-7](#)
 - entry timeout [5-12](#)
 - MAC learning, disabling [5-13](#)
 - resource management [6-18](#)
 - static entry [5-12](#)
- MAC learning, disabling [5-13](#)
- management interfaces
 - default settings [20-2, 21-2, 22-3](#)
- management IP address, transparent firewall [12-8](#)
- man-in-the-middle attack [5-6](#)
- mapped interface name [6-21](#)
- mask
 - reply, ICMP message [49-15](#)
 - request, ICMP message [49-15](#)
- Master Passphrase [13-8](#)
- message filtering [44-4](#)
- message list
 - filtering by [44-5](#)
- message-of-the-day banner [41-8](#)
- messages, logging
 - classes
 - about [44-4](#)
 - list of [44-4](#)
 - component descriptions [44-3](#)
 - filtering by message list [44-5](#)
 - format of [44-3](#)
 - message list, creating [44-14](#)
 - severity levels [44-3](#)
- messages classes [44-4](#)
- messages in EMBLEM format [44-15](#)
- metacharacters, regular expression [17-15](#)
- mgmt0 interfaces
 - default settings [20-2, 21-2, 22-3](#)
- MIBs [45-3](#)
- MIBs for SNMP [45-29](#)
- Microsoft Windows CA, supported [40-4](#)
- mobile redirection, ICMP message [49-16](#)
- mode
 - context [6-16](#)
 - firewall [5-1](#)
- modular policy framework
 - configuring flow-export actions for NetFlow [46-6](#)
- monitoring
 - failover [7-16](#)
 - OSPF [27-65](#)
 - resource management [6-30](#)
 - SNMP [45-1](#)
- monitoring logging [44-20](#)

monitoring NSEL [46-10](#)
 monitoring switch traffic, ASA 5505 [10-4](#)
 More prompt [48-5](#)
 MRoute pane
 description [30-5](#)
 MSFC
 overview [2-2](#)
 SVIs [2-5](#)
 MTU [11-12, 12-15](#)
 multicast traffic [5-5](#)
 multiple context mode
 logging [44-2](#)
 See security contexts
 multiple SVIs [2-5](#)

N

naming an interface
 other models [11-8, 12-11, 12-12](#)
 NAT
 disabling proxy ARP for global addresses [24-11](#)
 native VLAN support [10-10](#)
 neighbor reachable time [31-2](#)
 neighbor solicitation messages [31-2](#)
 neighbor advertisement messages [31-2](#)
 NetFlow
 overview [46-1](#)
 NetFlow collector
 configuring [46-5](#)
 NetFlow event
 matching to configured collectors [46-6](#)
 NetFlow event logging
 disabling [46-9](#)
 Network Activity test [7-17](#)
 No Payload Encryption [4-32](#)
 NSEL and syslog messages
 redundant messages [46-2](#)
 NSEL configuration examples [46-12](#)
 NSEL feature history [46-14](#)
 NSEL licensing requirements [46-4](#)
 NSEL runtime counters
 clearing [46-10](#)
 NT server
 configuring [34-15, 36-7, 37-3](#)

O

open ports [49-14](#)
 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 [44-5](#)
 output destinations [44-1, 44-7](#)
 e-mail address [44-1, 44-7](#)
 SNMP management station [44-1, 44-7](#)
 Telnet or SSH session [44-1, 44-7](#)
 outside, definition [1-14](#)
 oversubscribing resources [6-10](#)

P

packet
 capture [43-2](#)
 classifier [6-3](#)
 packet capture, enabling [43-3](#)
 packet flow
 routed firewall [5-14](#)
 transparent firewall [5-20](#)
 paging screen displays [48-5](#)
 parameter problem, ICMP message [49-15](#)
 passwords
 changing [13-3](#)
 recovery [13-14](#)
 security appliance [13-2](#)
 pause frames for flow control [9-26](#)
 PKI protocol [40-13](#)
 PoE [10-4](#)
 pools, address
 DHCP [14-5](#)
 port-forwarding
 enabling [11-7, 12-9](#)
 ports
 open on device [49-14](#)
 TCP and UDP [49-11](#)
 power over Ethernet [10-4](#)
 primary unit, failover [7-20](#)
 private networks [49-2](#)
 privileged EXEC mode
 accessing [3-4](#)
 privileged EXEC mode, accessing [3-1](#)

privileged mode
 accessing [3-1](#)
 prompt [48-2](#)
 Process Instances tab
 description [27-6](#)
 Product Authorization Key [4-35](#)
 prompts
 command [48-2](#)
 more [48-5](#)
 Properties tab [27-17](#)
 description [27-17](#)
 fields [27-17](#)
 protocol numbers and literal values [49-11](#)
 Protocol pane (PIM)
 description [30-10](#)
 proxy ARP, disabling [24-11](#)
 public key cryptography [40-2](#)

Q

question mark
 command string [48-4](#)
 help [48-4](#)
 queue, logging
 changing the size of [44-16](#)
 viewing statistics [44-20](#)

R

RADIUS
 attributes [34-3](#)
 configuring a server [34-15, 36-7, 37-3](#)
 support [34-1](#)
 rapid link failure detection [2-11](#)
 rate limit [44-20](#)
 redirect, ICMP message [49-15](#)
 redundant interface
 EtherChannel

- converting existing interfaces [9-16](#)
 - redundant interfaces
 - configuring [9-28](#)
 - failover [9-13](#)
 - MAC address [9-5](#)
 - setting the active interface [9-30](#)
 - Registration Authority description [40-2](#)
 - regular expression [17-14](#)
 - reloading
 - context [6-27](#)
 - security appliance [3-29](#)
 - Request Filter pane
 - description [30-12](#)
 - resetting the services module [2-12](#)
 - resource management
 - about [6-10](#)
 - assigning a context [6-22](#)
 - class [6-17](#)
 - configuring [6-8](#)
 - default class [6-9](#)
 - monitoring [6-30](#)
 - oversubscribing [6-10](#)
 - resource types [6-17](#)
 - unlimited [6-11](#)
 - resource usage [6-33](#)
 - revoked certificates [40-2](#)
 - RFCs for SNMP [45-29](#)
 - 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 [5-1](#)
 - setting [5-1](#)
 - route map
 - definition [26-1](#)
 - route maps
 - defining [26-4](#)
 - uses [26-1](#)
 - router
 - advertisement, ICMP message [49-15](#)
 - solicitation, ICMP message [49-15](#)
 - router advertisement messages [31-3](#)
 - router advertisement transmission interval [31-8](#)
 - router lifetime value [31-9](#)
 - routes
 - about default [25-4](#)
 - configuring default routes [25-4](#)
 - configuring IPv6 default [25-5](#)
 - configuring IPv6 static [25-5](#)
 - configuring static routes [25-3](#)
 - Route Summarization tab
 - description [27-6](#)
 - RSA
 - keys, generating [39-16, 39-18, 39-20, 39-22, 40-11, 41-4](#)
 - rules
 - ICMP [41-10](#)
 - running configuration
 - copying [42-25](#)
 - saving [3-25](#)
-
- ## S
- same security level communication
 - enabling [11-15, 12-18](#)
 - SDI
 - configuring [34-15, 36-7, 37-3](#)
 - secondary unit, failover [7-20](#)
 - Secure Copy
 - configure server [42-14](#)
 - security appliance
 - CLI [48-1](#)
 - connecting to [3-1](#)
 - managing licenses [4-1](#)

- managing the configuration [3-24](#)
- reloading [3-29](#)
- upgrading software [42-17](#)
- viewing files in Flash memory [42-12](#)
- security contexts
 - about [6-1](#)
 - adding [6-19](#)
 - admin context
 - about [6-2](#)
 - changing [6-26](#)
 - assigning to a resource class [6-22](#)
 - cascading [6-6](#)
 - changing between [6-24](#)
 - classifier [6-3](#)
 - command authorization [41-17](#)
 - configuration
 - URL, changing [6-26](#)
 - URL, setting [6-22](#)
 - logging in [6-7](#)
 - MAC addresses
 - automatically assigning [6-24](#)
 - classifying using [6-3](#)
 - managing [6-1, 6-25](#)
 - mapped interface name [6-21](#)
 - monitoring [6-28](#)
 - MSFC compatibility [2-3](#)
 - multiple mode, enabling [6-16](#)
 - nesting or cascading [6-7](#)
 - prompt [48-2](#)
 - reloading [6-27](#)
 - removing [6-25](#)
 - resource management [6-10](#)
 - resource usage [6-33](#)
 - saving all configurations [3-26](#)
 - unsupported features [6-14](#)
 - VLAN allocation [6-21](#)
- security level
 - about [11-1](#)
 - interface [11-9, 12-11, 12-13](#)
- security models for SNMP [45-16](#)
- segment size
 - maximum and minimum [11-10](#)
 - maximum and minimum, overview [9-8](#)
- sending messages to an e-mail address [44-11](#)
- sending messages to an SNMP server [44-12](#)
- sending messages to ASDM [44-12](#)
- sending messages to a specified output destination [44-17](#)
- sending messages to a syslog server [44-8](#)
- sending messages to a Telnet or SSH session [44-13](#)
- sending messages to the console port [44-12](#)
- sending messages to the internal log buffer [44-9](#)
- session management path [1-17](#)
- severity levels, of system log messages
 - changing [44-1](#)
 - filtering by [44-1](#)
 - list of [44-3](#)
- severity levels, of system messages
 - definition [44-3](#)
- shared license
 - backup server, configuring [4-39](#)
 - backup server, information [4-28](#)
 - client, configuring [4-39](#)
 - communication issues [4-28](#)
 - failover [4-29](#)
 - maximum clients [4-29](#)
 - monitoring [4-49](#)
 - server, configuring [4-37](#)
 - SSL messages [4-28](#)
- show command, filtering output [48-5](#)
- single mode
 - backing up configuration [6-16](#)
 - configuration [6-16](#)
 - enabling [6-16](#)
 - restoring [6-16](#)
- Smart Call Home monitoring [47-22](#)
- SNMP
 - about [45-1](#)
 - failover [45-17](#)

- management station [44-1, 44-7](#)
- prerequisites [45-17](#)
- SNMP configuration [45-18](#)
- SNMP groups [45-16](#)
- SNMP hosts [45-16](#)
- SNMP monitoring [45-26, 45-27](#)
- SNMP terminology [45-2](#)
- SNMP traps [45-3](#)
- SNMP users [45-16](#)
- SNMP Version 3 [45-15, 45-23](#)
- SNMP Versions 1 and 2c [45-22](#)
- source quench, ICMP message [49-15](#)
- SPAN [10-4](#)
- Spanning Tree Protocol, unsupported [10-8](#)
- SPAN session [2-6](#)
- speed, configuring [9-15, 10-6](#)
- SSH
 - authentication [41-21](#)
 - concurrent connections [41-2](#)
 - login [41-5](#)
 - password [13-2](#)
 - RSA key [41-4](#)
 - username [41-5](#)
- startup configuration
 - copying [42-25](#)
 - saving [3-25](#)
- Stateful Failover
 - about [7-13](#)
 - state information [7-13](#)
 - state link [7-4](#)
- stateful inspection [1-17](#)
- state information [7-13](#)
- state link [7-4](#)
- static ARP entry [5-10](#)
- static bridge entry [5-12](#)
- Static Group pane
 - description [30-7](#)
- static routes
 - configuring [25-3](#)
- stealth firewall
 - See* transparent firewall
- stuck-in-active [28-2](#)
- subcommand mode prompt [48-2](#)
- subinterfaces, adding [9-33](#)
- subnet masks
 - /bits [49-3](#)
 - about [49-2](#)
 - address range [49-4](#)
 - determining [49-3](#)
 - dotted decimal [49-3](#)
 - number of hosts [49-3](#)
- Summary Address pane
 - description [27-11](#)
- SVIs
 - configuring [2-10](#)
 - multiple [2-5](#)
 - overview [2-5](#)
- switch
 - assigning VLANs to module [2-7](#)
 - autostate messaging [2-11](#)
 - BPDU forwarding [2-11](#)
 - configuration [2-1](#)
 - failover compatibility with transparent firewall [2-11](#)
 - failover configuration [2-11](#)
 - trunk for failover [2-11](#)
 - verifying module installation [2-6](#)
- switched virtual interfaces
 - See* SVIs
- switch MAC address table [5-7](#)
- switch ports
 - access ports [10-7](#)
 - protected [10-8, 10-10](#)
 - SPAN [10-4](#)
 - trunk ports [10-9](#)
- SYN attacks, monitoring [6-34](#)
- SYN cookies [6-34](#)
- syntax formatting [48-3](#)
- syslogd server program [44-5](#)

- syslog messages
 - analyzing [44-2](#)
 - syslog messaging for SNMP [45-27](#)
 - syslog server
 - designating more than one as output destination [44-5](#)
 - EMBLEM format
 - configuring [44-15](#)
 - enabling [44-8, 44-15](#)
 - system configuration [6-2](#)
 - system log messages
 - classes [44-4](#)
 - classes of [44-4](#)
 - configuring in groups
 - by message list [44-5](#)
 - by severity level [44-1](#)
 - device ID, including [44-18](#)
 - disabling logging of [44-1](#)
 - filtering by message class [44-4](#)
 - managing in groups
 - by message class [44-17](#)
 - output destinations [44-1, 44-7](#)
 - syslog message server [44-7](#)
 - Telnet or SSH session [44-7](#)
 - severity levels
 - about [44-3](#)
 - changing the severity level of a message [44-1](#)
 - timestamp, including [44-19](#)
 - monitoring [6-34](#)
 - TCP MSS
 - overview [9-8](#)
 - Telnet
 - allowing management access [41-1](#)
 - authentication [41-21](#)
 - concurrent connections [41-2](#)
 - login [41-3](#)
 - password [13-2](#)
 - template timeout intervals
 - configuring for flow-export actions [46-7](#)
 - temporary license [4-24](#)
 - time exceeded, ICMP message [49-15](#)
 - time ranges, access lists [19-2](#)
 - timestamp, including in system log messages [44-19](#)
 - timestamp reply, ICMP message [49-15](#)
 - timestamp request, ICMP message [49-15](#)
 - traffic flow
 - routed firewall [5-14](#)
 - transparent firewall [5-20](#)
 - transparent firewall
 - about [5-2](#)
 - ARP inspection
 - about [5-6](#)
 - enabling [5-11](#)
 - static entry [5-10](#)
 - data flow [5-20](#)
 - guidelines [5-8](#)
 - H.323 guidelines [5-6](#)
 - HSRP [5-5](#)
 - MAC address timeout [5-12](#)
 - MAC learning, disabling [5-13](#)
 - management IP address [12-8](#)
 - multicast traffic [5-5](#)
 - static bridge entry [5-12](#)
 - unsupported features [5-9](#)
 - VRRP [5-5](#)
 - troubleshooting SNMP [45-24](#)
 - trunk, 802.1Q [9-33](#)
-
- ## T
- TACACS+
 - command authorization, configuring [41-33](#)
 - configuring a server [34-15, 36-7, 37-3](#)
 - TCP
 - connection limits per context [6-17](#)
 - maximum segment size [11-10](#)
 - maximum segment size, overview [9-8](#)
 - ports and literal values [49-11](#)
 - TCP Intercept

trunk ports [10-9](#)
 Trusted Flow Acceleration
 modes [5-8](#)
 trustpoint [40-3](#)

U

UDP
 connection limits per context [6-17](#)
 connection state information [1-17](#)
 ports and literal values [49-11](#)
 unprivileged mode
 accessing [3-4](#)
 unreachable, ICMP message [49-15](#)
 unreachable messages
 required for MTU discovery [41-10](#)
 upgrading
 IOS [2-3](#)
 URLs
 context configuration, changing [6-26](#)
 context configuration, setting [6-22](#)
 user EXEC mode
 accessing [3-1](#)
 prompt [48-2](#)
 username
 adding [33-4](#)
 encrypted [33-4](#)
 password [33-4](#)
 users
 SNMP [45-16](#)
 using clustering [44-5, 46-3](#)

V

VeriSign, configuring CAs example [40-4](#)
 viewing RMS [42-42](#)
 virtual firewalls
 See security contexts

Virtual Link
 description [27-28](#)
 virtual reassembly [1-15](#)
 VLANs [9-33](#)
 802.1Q trunk [9-33](#)
 allocating to a context [6-21](#)
 ASA 5505
 MAC addresses [10-4](#)
 maximum [10-2](#)
 assigning to FWSM [2-7](#)
 interfaces [2-7](#)
 mapped interface name [6-21](#)
 subinterfaces [9-33](#)
 VPN
 address range, subnets [49-4](#)
 VPN flex license [4-24](#)
 VRRP [5-5](#)

W

WCCP [16-1](#)
 web caching [16-1](#)

X

XOFF frames [9-26](#)

