



Cisco ASA Series General Operations ASDM Configuration Guide

Software Version 7.4

For the ASA 5506-X, ASA 5506H-X, ASA 5506W-X, ASA 5508-X, ASA 5512-X, ASA 5515-X, ASA 5516-X, ASA 5525-X, ASA 5545-X, ASA 5555-X, ASA 5585-X, ASA Services Module, and the Adaptive Security Virtual Appliance

First Published: March 23, 2015

Last Updated: April 7, 2015

Cisco Systems, Inc.

www.cisco.com

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

Text Part Number: N/A, Online only

<https://t.me/learningnets>

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

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

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

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

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

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

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

Cisco ASA Series General Operations ASDM Configuration Guide
Copyright © 2015 Cisco Systems, Inc. All rights reserved.



About This Guide

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

Document Objectives

The purpose of this guide is to help you configure general operations for the Cisco ASA series using the Adaptive Security Device Manager (ASDM). This guide does not cover every feature, but describes only the most common configuration scenarios.

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



Note

ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see [Cisco ASA Series Compatibility](#).

Related Documentation

For more information, see *Navigating the Cisco ASA Series Documentation* at <http://www.cisco.com/go/asadoocs>.

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/en/US/docs/general/whatsnew/whatsnew.html>.

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



PART 1

Getting Started with the ASA



Introduction to the Cisco ASA

First Published: March 23, 2015

Last Updated: April 7, 2015

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



Note

ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see [Cisco ASA Compatibility](#). See also [Special, Deprecated, and Legacy Services](#), page 1-17.

- [ASDM Requirements](#), page 1-1
- [Hardware and Software Compatibility](#), page 1-6
- [VPN Compatibility](#), page 1-6
- [New Features](#), page 1-6
- [Firewall Functional Overview](#), page 1-12
- [VPN Functional Overview](#), page 1-16
- [Security Context Overview](#), page 1-16
- [ASA Clustering Overview](#), page 1-17
- [Special, Deprecated, and Legacy Services](#), page 1-17

ASDM Requirements

- [ASDM Client Operating System and Browser Requirements](#), page 1-2
- [Java and Browser Compatibility](#), page 1-2

ASDM Client Operating System and Browser Requirements

The following table lists the supported and recommended client operating systems and Java for ASDM.

Table 1-1 *Operating System and Browser Requirements*

Operating System	Browser				Java SE Plug-in
	Internet Explorer	Firefox	Safari	Chrome	
Microsoft Windows (English and Japanese): 8 7 - Server 2008 - Server 2012	Yes	Yes	No support	Yes	7.0 or later
Apple OS X 10.4 and later	No support	Yes	Yes	Yes (64-bit version only)	7.0 or later
Red Hat Enterprise Linux 5 (GNOME or KDE): - Desktop - Desktop with Workstation	N/A	Yes	N/A	Yes	7.0 or later

Java and Browser Compatibility

The following table lists compatibility caveats for Java, ASDM, and browser compatibility.



Table 1-2 Java Caveats for ASDM Compatibility

Java Version	Conditions	Notes
7 update 51	ASDM Launcher requires trusted certificate	To continue using the Launcher, do one of the following: • Upgrade to Java 8 or downgrade Java to 7 update 45 or earlier. • Install a trusted certificate on the ASA from a known CA. • Install a self-signed certificate and register it with Java. See Install an Identity Certificate for ASDM. • Alternatively use Java Web Start. Note: ASDM 7.1(5) and earlier are not supported with Java 7 update 51. If you already upgraded Java, and can no longer launch ASDM in order to upgrade it to Version 7.2 or later, then you can either use the CLI to upgrade ASDM, or you can add a security exception in the Java Control Panel for each ASA you want to manage with ASDM. See the “Workaround” section at: http://java.com/en/download/help/java_blocked.xml After adding the security exception, launch the older ASDM and then upgrade to 7.2 or later.
	In rare cases, online help does not load when using Java Web Start	In rare cases, when launching online help, the browser window loads, but the content fails to appear. The browser reports an error: “Unable to connect”. Workaround: • Use the ASDM Launcher Or: • Clear the -Djava.net.preferIPv6Addresses=true parameter in Java Runtime Parameters: a. Launch the Java Control Panel. b. Click the Java tab. c. Click View. d. Clear this parameter: -Djava.net.preferIPv6Addresses=true e. Click OK, then Apply, then OK again.
7 update 45	ASDM shows a yellow warning about the missing Permissions attribute when using an untrusted certificate	Due to a bug in Java, if you do not have a trusted certificate installed on the ASA, you see a yellow warning about a missing Permissions attribute in the JAR manifest. It is safe to ignore this warning; ASDM 7.2 and later includes the Permissions attribute. To prevent the warning from appearing, install a trusted certificate (from a known CA); or generate a self-signed certificate on the ASA by choosing Configuration > Device Management > Certificates > Identity Certificates. Launch ASDM, and when the certificate warning is shown, check the Always trust connections to websites check box.

Table 1-2 Java Caveats for ASDM Compatibility (continued)

Java Version	Conditions	Notes
7	Requires strong encryption license (3DES/AES) on ASA	ASDM requires an SSL connection to the ASA. You can request a 3DES license from Cisco: 1. Go to www.cisco.com/go/license. 2. Click Continue to Product License Registration. 3. In the Licensing Portal, click Get Other Licenses next to the text field. 4. Choose IPS, Crypto, Other... from the drop-down list. 5. Type ASA in to the Search by Keyword field. 6. Select Cisco ASA 3DES/AES License in the Product list, and click Next. 7. Enter the serial number of the ASA, and follow the prompts to request a 3DES/AES license for the ASA.
All	• Self-signed certificate or an untrusted certificate • IPv6 • Firefox and Safari	When the ASA uses a self-signed certificate or an untrusted certificate, Firefox and Safari are unable to add security exceptions when browsing using HTTPS over IPv6. See https://bugzilla.mozilla.org/show_bug.cgi?id=633001. This caveat affects all SSL connections originating from Firefox or Safari to the ASA (including ASDM connections). To avoid this caveat, configure a proper certificate for the ASA that is issued by a trusted certificate authority.
	• SSL encryption on the ASA must include both RC4-MD5 and RC4-SHA1 <i>or</i> disable SSL false start in Chrome. • Chrome	If you change the SSL encryption on the ASA to exclude both RC4-MD5 and RC4-SHA1 algorithms (these algorithms are enabled by default), then Chrome cannot launch ASDM due to the Chrome “SSL false start” feature. We suggest re-enabling one of these algorithms (see the Configuration > Device Management > Advanced > SSL Settings pane); or you can disable SSL false start in Chrome using the --disable-ssl-false-start flag according to Run Chromium with flags.
	IE9 for servers	For Internet Explorer 9.0 for servers, the “Do not save encrypted pages to disk” option is enabled by default (See Tools > Internet Options > Advanced). This option causes the initial ASDM download to fail. Be sure to disable this option to allow ASDM to download.
	OS X	On OS X, you may be prompted to install Java the first time you run ASDM; follow the prompts as necessary. ASDM will launch after the installation completes.

Table 1-2 Java Caveats for ASDM Compatibility (continued)

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

Hardware and Software Compatibility

For a complete list of supported hardware and software, see [Cisco ASA Compatibility](#).

VPN Compatibility

See [Supported VPN Platforms, Cisco ASA Series](#).

New Features

Released: March 23, 2015

The following table lists the new features for ASA Version 9.4(1)/ASDM Version 7.4(1).

Table 1-3 *New Features for ASA Version 9.4(1)/ASDM Version 7.4(1)*

Feature	Description
Platform Features	
ASA 5506W-X, ASA 5506H-X, ASA 5508-X, ASA 5516-X	We introduced the following models: ASA 5506W-X with wireless access point, hardened ASA 5506H-X, ASA 5508-X, ASA 5516-X. We introduced the following command: hw-module module wlan recover image, hw-module module wlan recover image.
Certification Features	
Department of Defense Unified Capabilities Requirements (UCR) 2013 Certification	The ASA was updated to comply with the DoD UCR 2013 requirements. See the rows in this table for the following features that were added for this certification: • Periodic certificate authentication • Certificate expiration alerts • Enforcement of the basic constraints CA flag • ASDM Username From Certificate Configuration • IKEv2 invalid selectors notification configuration • IKEv2 pre-shared key in Hex

Table 1-3 **New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)**

Feature	Description
FIPS 140-2 Certification compliance updates	When you enable FIPS mode on the ASA, additional restrictions are put in place for the ASA to be FIPS 140-2 compliant. Restrictions include: • RSA and DH Key Size Restrictions—Only RSA and DH keys 2K (2048 bits) or larger are allowed. For DH, this means groups 1 (768 bit), 2 (1024 bit), and 5 (1536 bit) are not allowed. Note: The key size restrictions disable use of IKEv1 with FIPS. • Restrictions on the Hash Algorithm for Digital Signatures—Only SHA256 or better is allowed. • SSH Cipher Restrictions—Allowed ciphers: aes128-cbc or aes256-cbc. MACs: SHA1 To see the FIPS certification status for the ASA, see: http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140InProgress.pdf This PDF is updated weekly. See the Computer Security Division Computer Security Resource Center site for more information: http://csrc.nist.gov/groups/STM/cmvp/inprocess.html We modified the following command: fips enable
Firewall Features	
Improved SIP inspection performance on multiple core ASAs.	If you have multiple SIP signaling flows going through an ASA with multiple cores, SIP inspection performance has been improved. However, you will not see improved performance if you are using a TLS, phone, or IME proxy. We did not modify any screens.
SIP inspection support for Phone Proxy and UC-IME Proxy was removed.	You can no longer use Phone Proxy or UC-IME Proxy when configuring SIP inspection. Use TLS Proxy to inspect encrypted traffic. We removed Phone Proxy and UC-IME Proxy from the Select SIP Inspect Map service policy dialog box.
DCERPC inspection support for ISystemMapper UUID message RemoteGetClassObject opnum3.	The ASA started supporting non-EPM DCERPC messages in release 8.3, supporting the ISystemMapper UUID message RemoteCreateInstance opnum4. This change extends support to the RemoteGetClassObject opnum3 message. We did not modify any screens.
Unlimited SNMP server trap hosts per context	The ASA supports an unlimited number of SNMP server trap hosts per context. The show snmp-server host command output displays only the active hosts that are polling the ASA, as well as the statically configured hosts. We did not modify any screens.
VXLAN packet inspection	The ASA can inspect the VXLAN header to enforce compliance with the standard format. We modified the following screen: Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > Protocol Inspection

Table 1-3 *New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)*

Feature	Description
DHCP monitoring for IPv6	You can now monitor DHCP statistics and DHCP bindings for IPv6. We introduced the following screens: Monitoring > Interfaces > DHCP > IPV6 DHCP Statistics Monitoring > Interfaces > DHCP > IPV6 DHCP Binding.
High Availability Features	
Blocking syslog generation on a standby ASA	You can now block specific syslogs from being generated on a standby unit. We did not modify any screens.
Enable and disable ASA cluster health monitoring per interface	You can now enable or disable health monitoring per interface. Health monitoring is enabled by default on all port-channel, redundant, and single physical interfaces. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored. You might want to disable health monitoring of non-essential interfaces, for example, the management interface. We introduced the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Interface Health Monitoring
ASA clustering support for DHCP relay	You can now configure DHCP relay on the ASA cluster. Client DHCP requests are load-balanced to the cluster members using a hash of the client MAC address. DHCP client and server functions are still not supported. We did not modify any screens.
SIP inspection support in ASA clustering	You can now configure SIP inspection on the ASA cluster. A control flow can be created on any unit (due to load balancing), but its child data flows must reside on the same unit. TLS Proxy configuration is not supported. We did not modify any screens.
Routing Features	
Policy Based Routing	Policy Based Routing (PBR) is a mechanism by which traffic is routed through specific paths with a specified QoS using ACLs. ACLs let traffic be classified based on the content of the packet's Layer 3 and Layer 4 headers. This solution lets administrators provide QoS to differentiated traffic, distribute interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths, and allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections. We introduced or modified the following screens: Configuration > Device Setup > Routing > Route Maps > Policy Based Routing Configuration > Device Setup > Routing > Interface Settings > Interfaces.
Interface Features	
VXLAN support	VXLAN support was added, including VXLAN tunnel endpoint (VTEP) support. You can define one VTEP source interface per ASA or security context. We introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces > Add > VNI Interface Configuration > Device Setup > Interface Settings > VXLAN
Monitoring Features	

Table 1-3 ***New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)***

Feature	Description
Memory tracking for the EEM	We have added a new debugging feature to log memory allocations and memory usage, and to respond to memory logging wrap events. We modified the following screen: Configuration > Device Management > Advanced > Embedded Event Manager > Add Event Manager Applet > Add Event Manager Applet Event.
Troubleshooting crashes	The show tech-support command output and show crashinfo command output includes the most recent 50 lines of generated syslogs. Note that you must enable the logging buffer command to enable these results to appear.
Remote Access Features	
Support for ECDHE-ECDSA ciphers	TLSv1.2 added support for the following ciphers: • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-GCM-SHA384 • AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA256 Note: ECDSA and DHE ciphers are the highest priority. We modified the following screen: Configuration > Remote Access VPN > Advanced > SSL Settings.

Table 1-3 **New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)**

Feature	Description
Clientless SSL VPN session cookie access restriction	You can now prevent a Clientless SSL VPN session cookie from being accessed by a third party through a client-side script such as Javascript. Note: Use this feature only if Cisco TAC advises you to do so. Enabling this command presents a security risk because the following Clientless SSL VPN features will not work without any warning. • Java plug-ins • Java rewriter • Port forwarding • File browser • Sharepoint features that require desktop applications (for example, MS Office applications) • AnyConnect Web launch • Citrix Receiver, XenDesktop, and Xenon • Other non-browser-based and browser plugin-based applications We introduced the following screen: Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > HTTP Cookie. <i>This feature is also in 9.2(3).</i>
Virtual desktop access control using security group tagging	The ASA now supports security group tagging-based policy control for Clientless SSL remote access to internal applications and websites. This feature uses Citrix's virtual desktop infrastructure (VDI) with XenDesktop as the delivery controller and the ASA's content transformation engine. See the following Citrix product documentation for more information: • Policies for XenDesktop and XenApp: http://support.citrix.com/proddocs/topic/infocenter/ic-how-to-use.html • Managing policies in XenDesktop 7: http://support.citrix.com/proddocs/topic/xendesktop-7/cds-policies-wrapper-rho.html • Using group policy editor for XenDesktop 7 policies: http://support.citrix.com/proddocs/topic/xendesktop-7/cds-policies-use-gpms.html
OWA 2013 feature support has been added for Clientless SSL VPN	Clientless SSL VPN supports the new features in OWA 2013 except for the following: • Support for tablets and smartphones • Offline mode • Active Directory Federation Services (AD FS) 2.0. The ASA and AD FS 2.0 can't negotiate encryption protocols. We did not modify any screens.

Table 1-3 **New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)**

Feature	Description
Citrix XenDesktop 7.5 and StoreFront 2.5 support has been added for Clientless SSL VPN	Clientless SSL VPN supports the access of XenDesktop 7.5 and StoreFront 2.5. See http://support.citrix.com/proddocs/topic/xenapp-xendesktop-75/cds-75-about-whats-new.html for the full list of XenDesktop 7.5 features, and for more details. See http://support.citrix.com/proddocs/topic/dws-storefront-25/dws-about.html for the full list of StoreFront 2.5 features, and for more details. We did not modify any screens.
Periodic certificate authentication	When you enable periodic certificate authentication, the ASA stores certificate chains received from VPN clients and re-authenticates them periodically. We modified the following screens: Configuration > Device Management > Certificate Management > Identity Certificates Configuration > Device Management > Certificate Management > CA Certificates
Certificate expiration alerts	The ASA checks all CA and ID certificates in the trust points for expiration once every 24 hours. If a certificate is nearing expiration, a syslog will be issued as an alert. You can configure the reminder and recurrence intervals. By default, reminders will start at 60 days prior to expiration and recur every 7 days. We modified the following screens: Configuration > Device Management > Certificate Management > Identity Certificates Configuration > Device Management > Certificate Management > CA Certificates
Enforcement of the basic constraints CA flag	Certificates without the CA flag now cannot be installed on the ASA as CA certificates by default. The basic constraints extension identifies whether the subject of the certificate is a CA and the maximum depth of valid certification paths that include this certificate. You can configure the ASA to allow installation of these certificates if desired. We modified the following screens: Configuration > Device Management > Certificate Management > CA Certificates
IKEv2 invalid selectors notification configuration	Currently, if the ASA receives an inbound packet on an SA, and the packet's header fields are not consistent with the selectors for the SA, then the ASA discards the packet. You can now enable or disable sending an IKEv2 notification to the peer. Sending this notification is disabled by default. Note: This feature is supported with AnyConnect 3.1.06060 and later.
IKEv2 pre-shared key in Hex	You can now configure the IKEv2 pre-shared keys in hex. We modified the following screen: Configuration > Site-to-Site VPN > Connection Profiles
Administrative Features	

Table 1-3 **New Features for ASA Version 9.4(1)/ASDM Version 7.4(1) (continued)**

Feature	Description
ASDM Username From Certificate Configuration	This feature introduces the ability to authorize ASDM users by extracting the username from the certificate along with using the username provided by the user. We introduced the following screen: Configuration > Device Management > Management Access > HTTP Certificate Rule We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization
terminal interactive command to enable or disable help when you enter ? at the CLI	Normally, when you enter ? at the ASA CLI, you see command help. To be able to enter ? as text within a command (for example, to include a ? as part of a URL), you can disable interactive help using the no terminal interactive command.
REST API Version 1.1	We added support for the REST API Version 1.1.

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.

- [Security Policy Overview, page 1-12](#)
- [Firewall Mode Overview, page 1-14](#)
- [Stateful Inspection Overview, page 1-15](#)

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.

- [Permitting or Denying Traffic with Access Rules, page 1-13](#)
- [Applying NAT, page 1-13](#)

- [Protecting from IP Fragments, page 1-13](#)
- [Applying HTTP, HTTPS, or FTP Filtering, page 1-13](#)
- [Applying Application Inspection, page 1-13](#)
- [Sending Traffic to Supported Hardware or Software Modules, page 1-14](#)
- [Applying QoS Policies, page 1-14](#)
- [Applying Connection Limits and TCP Normalization, page 1-14](#)
- [Enabling Threat Detection, page 1-14](#)

Permitting or Denying Traffic with Access Rules

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

Applying NAT

Some of the benefits of NAT include the following:

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

Protecting from IP Fragments

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

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.

You can configure Cloud Web Security on the ASA, or install an ASA module that provides URL and other filtering services, such as ASA CX or ASA FirePOWER. You can also use the ASA in conjunction with an external product such as the Cisco Web Security Appliance (WSA).

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 Supported Hardware or Software Modules

Some ASA models allow you to configure software modules, or to insert hardware modules into the chassis, to provide advanced services. These modules provide additional traffic inspection and can block traffic based on your configured policies. You can send traffic to these modules to take advantage of these advanced services.

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.

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.

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; however, some features are not supported. See the feature chapters for more information.

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.

Special, Deprecated, and Legacy Services

For some services, documentation is located outside of the main configuration guides and online help. For a complete list of guides, see [Navigating the Cisco ASA Series Documentation](#).

- [Special Services Guides, page 1-17](#)
- [Deprecated Services, page 1-17](#)
- [Legacy Services Guide, page 1-17](#)

Special Services Guides

Special services allow the ASA to interoperate with other Cisco products; for example, by providing a security proxy for phone services (Unified Communications), or by providing Botnet traffic filtering in conjunction with the dynamic database from the Cisco update server, or by providing WCCP services for the Cisco Web Security Appliance. Some of these special services are covered in separate guides.

Deprecated Services

For deprecated features, see the configuration guide for your ASA version. Similarly, for redesigned features such as NAT between Version 8.2 and 8.3 or transparent mode interfaces between Version 8.3 and 8.4, refer to the configuration guide for your version. Although ASDM is backwards compatible with previous ASA releases, the configuration guide and online help only cover the latest release.

Legacy Services Guide

Legacy services are still supported on the ASA, however there may be better alternative services that you can use instead. Legacy services are covered in a separate guide.



Getting Started

This chapter describes how to get started with your Cisco ASA.

- [Access the Console for Command-Line Interface, page 2-1](#)
- [Configure ASDM Access, page 2-7](#)
- [Start ASDM, page 2-13](#)
- [Customize ASDM Operation, page 2-14](#)
- [Factory Default Configurations, page 2-15](#)
- [Get Started with the Configuration, page 2-21](#)
- [Use the Command Line Interface Tool in ASDM, page 2-22](#)
- [Apply Configuration Changes to Connections, page 2-23](#)

Access the Console for Command-Line Interface

In some cases, you may need to use the CLI to configure basic settings for ASDM access.

For initial configuration, access the CLI directly from the console port. Later, you can configure remote access using Telnet or SSH according to [Chapter 34, “Management Access.”](#) If your system is already in multiple context mode, then accessing the console port places you in the system execution space.



Note

For ASAv console access, see the ASAv quick start guide.

- [Access the Appliance Console, page 2-1](#)
- [Access the ASA Services Module Console, page 2-2](#)
- [Access the Software Module Console, page 2-6](#)
- [Access the ASA 5506W-X Wireless Access Point Console, page 2-7](#)

Access the Appliance Console

Follow these steps to access the appliance console.

Procedure

-
- Step 1** Connect a computer 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 [Set the Hostname, Domain Name, and the Enable and Telnet Passwords, page 18-1](#) 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.
-

Access the ASA Services Module Console

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

- [About Connection Methods, page 2-3](#)
- [Log Into the ASA Services Module, page 2-4](#)
- [Log Out of a Console Session, page 2-5](#)
- [Kill an Active Console Connection, page 2-5](#)

- [Log Out of a Telnet Session, page 2-6](#)

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 software, 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.

- 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.

Log 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.

Later, you can configure remote access directly to the ASASM using Telnet or SSH.

Procedure

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

- Available for initial access—From the switch CLI, enter this command to gain console access to the ASASM:

```
service-module session [switch {1 | 2}] slot number
```

Example:

```
Router# service-module session slot 3
ciscoasa>
```

For a switch in a VSS, enter the **switch** argument.

To view the module slot numbers, enter the **show module** command at the switch prompt.

You access user EXEC mode.

- Available after you configure a login password—From the switch CLI, enter this command to Telnet to the ASASM over the backplane:

```
session [switch {1 | 2}] slot number processor 1
```

You are prompted for the login password:

```
ciscoasa passwd:
```

Example:

```
Router# session slot 3 processor 1
ciscoasa passwd: cisco
ciscoasa>
```

For a switch in a VSS, enter the **switch** argument.

The **session slot processor 0** command, which is supported on other services modules, is not supported on the ASASM; the ASASM does not have a processor 0.

To view the module slot numbers, enter the **show module** command at the switch prompt.

Enter the login password to the ASASM. Set the password using the **passwd** command. There is no default password.

You access user EXEC mode.

Step 2 Access privileged EXEC mode, which is the highest privilege level:

```
enable
```

Example:

```
ciscoasa> enable
Password:
ciscoasa#
```


Enter the enable password at the prompt. By default, the password is blank.

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

Step 3 Access global configuration mode:

```
configure terminal
```

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

Related Topics

- [Guidelines for Management Access, page 34-1.](#)
- [Set the Hostname, Domain Name, and the Enable and Telnet Passwords, page 18-1](#)

Log Out of a Console Session

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 [Kill an Active Console Connection, page 2-5](#).

Procedure

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**.

Kill 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.

Procedure

- 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
```

Log Out of a Telnet Session

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

Procedure

- 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.

Access the Software Module Console

If you have a software module installed, such as the ASA FirePOWER module on the ASA 5506-X, you can session to the module console.

**Note**

You cannot access the *hardware* module CLI over the ASA backplane using the **session** command.

Procedure

Step 1 From the ASA CLI, session to the module:

```
session {sfr | cxsc | ips} console
```

Example:

```
ciscoasa# session sfr console
Opening console session with module sfr.
Connected to module sfr. Escape character sequence is 'CTRL-^X'.

Cisco ASA SFR Boot Image 5.3.1
asasfr login: admin
Password: Admin123
```

Access the ASA 5506W-X Wireless Access Point Console

To access the wireless access point console, perform the following steps.

Procedure

Step 1 From the ASA CLI, session to the access point:

```
session wlan console
```

Example:

```
ciscoasa# session wlan console
opening console session with module wlan
connected to module wlan. Escape character sequence is 'CTRL-^X'

ap>
```

Step 2 See the [Cisco IOS Configuration Guide for Autonomous Aironet Access Points](#) for information about the access point CLI.

Configure ASDM Access

This section describes how to access ASDM with a default configuration and how to configure access if you do not have a default configuration.

- [Use the Factory Default Configuration for ASDM Access \(Appliances, ASAv\), page 2-8](#)
- [Customize ASDM Access for Appliances and the ASAv, page 2-8](#)
- [Configure ASDM Access for the ASA Services Module, page 2-10](#)

Use the Factory Default Configuration for ASDM Access (Appliances, ASAv)

With a factory default configuration, ASDM connectivity is pre-configured with default network settings.

Procedure

- Step 1** Connect to ASDM using the following interface and network settings:
- The management interface depends on your model:
 - ASA 5506-X, ASA 5508-X, and ASA 5516-X—The interface to which you connect to ASDM is GigabitEthernet 1/2.
 - ASA 5512-X and higher—The interface to which you connect to ASDM is Management 0/0.
 - ASAv—The interface to which you connect to ASDM is Management 0/0.
 - The default management address is:
 - ASA appliances—192.168.1.1.
 - ASAv—You set the management interface IP address during deployment.
 - The clients allowed to access ASDM:
 - ASA appliances—Clients must be on the 192.168.1.0/24 network. The default configuration enables DHCP so that your management station can be assigned an IP address in this range.
 - ASAv—You set the management client IP address during deployment. The ASAv does not act as the DHCP server for connected clients.



Note

If you change to multiple context mode, you can access ASDM from the admin context using the network settings above.

Related Topics

- [Factory Default Configurations, page 2-15](#)
- [Enable or Disable Multiple Context Mode, page 8-15](#)
- [Start ASDM, page 2-13](#)

Customize ASDM Access for Appliances and the ASAv

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

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

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. 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.

**Note**

For the ASAv, you can configure transparent mode when you deploy, so this procedure is primarily useful after you deploy if you need to clear your configuration, for example.

Procedure

Step 1 Access the CLI at the console port.

Step 2 (Optional) Enable transparent firewall mode:

This command clears your configuration.

```
firewall transparent
```

Step 3 Configure the management interface:

```
interface interface_id  
  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
```

The **security-level** is a number between 1 and 100, where 100 is the most secure.

Step 4 (For directly-connected management hosts) Set the DHCP pool for the management network:

```
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
```

Make sure you do not include the interface address in the range.

Step 5 (For remote management hosts) Configure a route to the 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 1
```

Step 6 Enable the HTTP server for ASDM:

```
http server enable
```

Step 7 Allow the management host(s) to access ASDM:

```
http ip_address mask interface_name
```

Example:

```
ciscoasa(config)# http 192.168.1.0 255.255.255.0 management
```

Step 8 Save the configuration:

```
write memory
```

Step 9 (Optional) Set the mode to multiple mode:

```
mode multiple
```

When prompted, confirm that you want to convert the existing configuration to be the admin context. You are then prompted to reload the ASA.

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
```

Related Topics

- [Restore the Factory Default Configuration, page 2-16](#)
- [Set the Firewall Mode \(Single Mode\), page 6-9](#)
- [Access the Appliance Console, page 2-1](#)
- [Start ASDM, page 2-13](#)
- [Chapter 8, “Multiple Context Mode.”](#)

Configure 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.

Before You Begin

Assign a VLAN interface to the ASASM according to ASASM quick start guide.

Procedure

Step 1 Connect to the ASASM and access global configuration mode.

Step 2 (Optional) Enable transparent firewall mode:

```
firewall transparent
```

This command clears your configuration.

Step 3 Do one of the following to configure a management interface, depending on your mode:

- Routed mode—Configure an interface in routed mode:

```
interface vlan number
  ip address ip_address [mask]
  nameif name
  security-level level
```

Example:

```
ciscoasa(config)# interface vlan 1
ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
```

The **security-level** is a number between 1 and 100, where 100 is the most secure.

- Transparent mode—Configure a bridge virtual interface and assigns a management VLAN to the bridge group:

```
interface bvi number
  ip address ip_address [mask]

interface vlan number
  bridge-group bvi_number
  nameif name
  security-level level
```

Example:

```
ciscoasa(config)# interface bvi 1
ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0

ciscoasa(config)# interface vlan 1
ciscoasa(config-if)# bridge-group 1
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
```

The **security-level** is a number between 1 and 100, where 100 is the most secure.

Step 4 (For directly-connected management hosts) Enable DHCP for the management host on the management interface network:

```
dhcpd address ip_address-ip_address
interface_name
dhcpd enable interface_name
```

Example:

```
ciscoasa(config)# dhcpd address 192.168.1.2-192.168.1.254 inside
ciscoasa(config)# dhcpd enable inside
```

Make sure you do not include the management address in the range.

Step 5 (For remote management hosts) Configure a route to the 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
```

Step 6 Enable the HTTP server for ASDM:

```
http server enable
```

Step 7 Allow the management host to access ASDM:

```
http ip_address mask interface_name
```

Example:

```
ciscoasa(config)# http 192.168.1.0 255.255.255.0 management
```

Step 8 Save the configuration:

```
write memory
```

Step 9 (Optional) Set the mode to multiple mode:

```
mode multiple
```

When prompted, confirm that you want to convert the existing configuration to be the admin context. You are then prompted to reload the ASDM.

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
```

Related Topics

- [Access the ASA Services Module Console, page 2-2](#)
- [Chapter 8, “Multiple Context Mode.”](#)
- [Set the Firewall Mode \(Single Mode\), page 6-9](#)

Start 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 shortcut to your computer; however you need separate shortcuts for each ASA IP address.

Within ASDM, you can choose a different ASA IP address to manage; the difference between the Launcher and Java Web Start 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.

Procedure

Step 1 On the computer that you specified as the ASDM client, enter the following URL:

`https://asa_ip_address/admin`

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. Leave the username and password fields empty (for a new installation), and click **OK**. With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. Note: If you enabled HTTPS authentication, enter your username and associated password.
- c. Save the installer to your computer, and then start the installer. The ASDM-IDM Launcher opens automatically after installation is complete.
- d. Enter the management IP address, leave the username and password blank (for a new installation), and then click **OK**. Note: If you enabled HTTPS authentication, enter your username and associated password.

Step 3 To use Java Web Start:

- a. Click **Run ASDM** or **Run Startup Wizard**.
- b. Save the shortcut to your computer when prompted. You can optionally open it instead of saving it.
- c. Start Java Web Start from the shortcut.
- d. Accept any certificates according to the dialog boxes that appear. The Cisco ASDM-IDM Launcher appears.

- e. Leave the username and password blank (for a new installation), and then click **OK**. Note: If you enabled HTTPS authentication, enter your username and associated password.
-

Customize ASDM Operation

You can install an identity certificate to successfully launch ASDM as well as increase the ASDM heap memory so it can handle larger configurations.

- [Install an Identity Certificate for ASDM, page 2-14](#)
- [Increase the ASDM Configuration Memory, page 2-14](#)

Install an Identity Certificate for ASDM

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

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

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

Increase the ASDM Configuration Memory

ASDM supports a maximum configuration size of 512 KB. If you exceed this amount you may experience performance issues. For example, when you load the configuration, the status dialog box shows the percentage of the configuration that is complete, yet with large configurations it stops incrementing and appears to suspend operation, even though ASDM might still be processing the configuration. If this situation occurs, we recommend that you consider increasing the ASDM system heap memory.

- [Increase the ASDM Configuration Memory in Windows, page 2-14](#)
- [Increase the ASDM Configuration Memory in Mac OS, page 2-15](#)

Increase the ASDM Configuration Memory in Windows

To increase the ASDM heap memory size, edit the **run.bat** file by performing the following procedure.

Procedure

1. Go to the ASDM installation directory, for example C:\Program Files (x86)\Cisco Systems\ASDM.
2. Edit the **run.bat** file with any text editor.
3. In the line that starts with “start javaw.exe”, change the argument prefixed with “-Xmx” to specify your desired heap size. For example, change it to -Xmx768M for 768 MB or -Xmx1G for 1 GB.
4. Save the **run.bat** file.

Increase the ASDM Configuration Memory in Mac OS

To increase the ASDM heap memory size, edit the **Info.plist** file by performing the following procedure.

Procedure


1. Right-click the **Cisco ASDM-IDM** icon, and choose **Show Package Contents**.
2. In the **Contents** folder, double-click the **Info.plist** file. If you have Developer tools installed, it opens in the **Property List Editor**. Otherwise, it opens in **TextEdit**.
3. Under **Java > VMOptions**, change the string prefixed with “-Xmx” to specify your desired heap size. For example, change it to -Xmx768M for 768 MB or -Xmx1G for 1 GB.

```
<key>CFBundleIconFile</key>
<string>asdm32.icns</string>

<key>VMOptions</key>
<string>-Xms64m -Xmx512m</string>
```

```
<key>CFBundleDocumentTypes</key>
<array>
```

4. If this file is locked, you see an error such as the following:



5. Click **Unlock** and save the file.

If you do not see the **Unlock** dialog box, exit the editor, right-click the **Cisco ASDM-IDM** icon, choose **Copy Cisco ASDM-IDM**, and paste it to a location where you have write permissions, such as the Desktop. Then change the heap size from this copy.

Factory Default Configurations

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

- ASA appliances—The factory default configuration configures an interface for management so that you can connect to it using ASDM, with which you can then complete your configuration.
- ASAv—Depending on your hypervisor, as part of deployment, the deployment configuration (the initial virtual deployment settings) configures an interface for management so that you can connect to it using ASDM, with which you can then complete your configuration. You can also configure failover IP addresses. You can also apply a “factory default” configuration if desired.

- ASASM—No default configuration. See [Access the ASA Services Module Console, page 2-2](#) to start configuration.

For appliances, the factory default configuration is available only for routed firewall mode and single context mode. For the ASAv, you can choose transparent or routed mode at deployment.

**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.

- [Restore the Factory Default Configuration, page 2-16](#)
- [Restore the ASAv Deployment Configuration, page 2-17](#)
- [ASA 5506-X, 5508-X, and 5516-X Default Configuration, page 2-18](#)
- [ASA 5512-X, 5515-X, 5525-X and Above Default Configuration, page 2-19](#)
- [ASAv Deployment Configuration, page 2-20](#)

Restore the Factory Default Configuration

This section describes how to restore the factory default configuration. Both CLI and ASDM procedures are provided. For the ASAv, this procedure erases the deployment configuration and applies the same factory default configuration as for the ASA appliances.

**Note**

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

Before You Begin

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.

Procedure

Step 1 Restore the factory default configuration:

```
configure factory-default [ip_address [mask]]
```

Example:

```
ciscoasa(config)# configure factory-default 10.1.1.1 255.255.255.0
```

If you specify the *ip_address*, then you set the inside or management interface IP address, depending on your model, instead of using the default IP address of 192.168.1.1. The **http** command uses the subnet you specify. Similarly, the **dhcpd address** command range consists of addresses within the subnet that you specify.

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

Step 2 Save the default configuration to flash memory:

```
write 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.

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

The **Reset Device to the Default Configuration** dialog box appears.

Step 2 (Optional) Enter the **Management IP address** of the management interface, instead of using the default address, 192.168.1.1.

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

Step 4 Click **OK**.

A confirmation dialog box appears.



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

Step 5 Click **Yes**.

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

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

Restore the ASAv Deployment Configuration

This section describes how to restore the ASAv deployment (Day 0) configuration.

Procedure

Step 1 For failover, power off the standby unit.

To prevent the standby unit from becoming active, you must power it off. If you leave it on, when you erase the active unit configuration, then the standby unit becomes active. When the former active unit reloads and reconnects over the failover link, the old configuration will sync from the new active unit, wiping out the deployment configuration you wanted.

- Step 2** Restore the deployment configuration after you reload. For failover, enter this command on the active unit:

```
write erase
```



Note

The ASAv boots the current running image, so you are not reverted to the original boot image. To use the original boot image, see the **boot image** command.

Do not save the configuration.

- Step 3** Reload the ASAv and load the deployment configuration:

```
reload
```

- Step 4** For failover, power on the standby unit.

After the active unit reloads, power on the standby unit. The deployment configuration will sync to the standby unit.

ASA 5506-X, 5508-X, and 5516-X Default Configuration

The default factory configuration for the ASA 5506-X series, 5508-X, and 5516-X configures the following:

- inside --> outside traffic flow—GigabitEthernet 1/1 (outside), GigabitEthernet 1/2 (inside)
- outside IP address from DHCP, inside IP address—192.168.1.1
- (ASA 5506W-X) wifi <--> inside, wifi --> outside traffic flow—GigabitEthernet 1/9 (wifi)
- (ASA 5506W-X) wifi IP address—192.168.10.1
- DHCP for clients on inside and wifi. The access point itself and all its clients use the ASA as the DHCP server.
- Management 1/1 interface is Up, but otherwise unconfigured. The ASA FirePOWER module can then use this interface to access the ASA inside network and use the inside interface as the gateway to the Internet.
- ASDM access—inside and wifi hosts allowed.
- NAT—Interface PAT for all traffic from inside, wfi, and management to outside.

The configuration consists of the following commands:

```
interface Management1/1
  management-only
  no nameif
  no security-level
  no ip address
  no shutdown
interface GigabitEthernet1/1
  nameif outside
  security-level 0
```

```
ip address dhcp setroute
no shutdown
interface GigabitEthernet1/2
  nameif inside
  security-level 100
  ip address 192.168.1.1 255.255.255.0
  no shutdown
object network obj_any
  subnet 0.0.0.0 0.0.0.0
  nat (any,outside) dynamic interface
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd auto_config outside
dhcpd address 192.168.1.5-192.168.1.254 inside
dhcpd enable inside
logging asdm informational
```

For the ASA 5506W-X, the following commands are also included:

```
same-security-traffic permit inter-interface
interface GigabitEthernet 1/9
  security-level 100
  nameif wifi
  ip address 192.168.10.1 255.255.255.0
  no shutdown
http 192.168.10.0 255.255.255.0 wifi
dhcpd address 192.168.10.2-192.168.10.254 wifi
dhcpd enable wifi
```

ASA 5512-X, 5515-X, 5525-X and Above Default Configuration

The default factory configuration for the ASA 5512-X, 5515-X, 5525-X and above 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 that a computer 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
```

ASAv Deployment Configuration

When you deploy the ASAv, you can pre-set many parameters that let you connect to the Management 0/0 interface using ASDM. A typical configuration includes the following settings:

- Routed or Transparent firewall mode
- Management 0/0 interface:
 - Named “management”
 - IP address or DHCP
 - Security level 0
 - Management-only
- Static route for the management host IP address (if it is not on the management subnet)
- HTTP server enabled or disabled
- HTTP access for the management host IP address
- (Optional) Failover link IP addresses for GigabitEthernet 0/8, and the Management 0/0 standby IP address
- DNS server
- Smart licensing ID token
- Smart licensing Throughput Level and Standard Feature Tier
- (Optional) Smart Call Home HTTP Proxy URL and port
- (Optional) SSH management settings:
 - Client IP addresses
 - Local username and password
 - Authentication required for SSH using the LOCAL database
- (Optional) REST API enabled or disabled



Note

To successfully register the ASAv with the Cisco Licensing Authority, the ASAv requires Internet access. You might need to perform additional configuration after deployment to achieve Internet access and successful license registration.

See the following sample configuration for a standalone unit:

```
interface Management0/0
  nameif management
  security-level 0
  ip address ip_address
  management-only
  no shutdown
http server enable
http management_host_IP mask management
route management management_host_IP mask gateway_ip 1
dns server-group DefaultDNS
  name-server ip_address
call-home
  http-proxy ip_address port port
license smart
  feature tier standard
  throughput level {100M | 1G | 2G}
```



```
license smart register idtoken id_token
aaa authentication ssh console LOCAL
username username password password
ssh source_IP_address mask management
rest-api image boot:/path
rest-api agent
```

See the following sample configuration for a primary unit in a failover pair:

```
interface Management0/0
  nameif management
  security-level 0
  ip address ip_address standby standby_ip
  management-only
  no shutdown
route management management_host_IP mask gateway_ip 1
http server enable
http management_host_IP mask management
dns server-group DefaultDNS
  name-server ip_address
call-home
  http-proxy ip_address port port
license smart
  feature tier standard
  throughput level {100M | 1G | 2G}
license smart register idtoken id_token
aaa authentication ssh console LOCAL
username username password password
ssh source_IP_address mask management
rest-api image boot:/path
rest-api agent
failover
failover lan unit primary
failover lan interface fover gigabitethernet0/8
failover link fover gigabitethernet0/8
failover interface ip fover primary_ip mask standby standby_ip
```

Get Started with the Configuration

To configure and monitor the ASA, perform the following steps.



Note

ASDM supports up to a maximum of a 512 KB configuration. If you exceed this amount, you may experience performance issues. See [Increase the ASDM Configuration Memory, page 2-14](#).

Procedure

- Step 1** For initial configuration using the Startup Wizard, choose **Wizards > Startup Wizard**.
- Step 2** To use the IPsec [VPN Wizard](#) to configure IPsec VPN connections, choose **Wizards > IPsec VPN Wizard** and complete each screen that appears.
- Step 3** To use the SSL [VPN Wizard](#) to configure SSL VPN connections, choose **Wizards > SSL VPN Wizard** and complete each screen that appears.
- Step 4** To configure high availability and scalability settings, choose **Wizards > High Availability and Scalability Wizard**.

- Step 5** To use the Packet Capture Wizard to configure packet capture, choose **Wizards > Packet Capture Wizard**.
- Step 6** To display different colors and styles available in the ASDM GUI, choose **View > Office Look and Feel**.
- Step 7** To configure features, click the **Configuration** button on the toolbar and then click one of the feature buttons to display the associated configuration pane.



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

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

Use the Command Line Interface Tool in ASDM

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

- [Use the Command Line Interface Tool, page 2-22](#)
- [Show Commands Ignored by ASDM on the Device, page 2-23](#)

Use the Command Line Interface Tool

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

The commands you can enter with the CLI tool depend on your user privileges. Review your privilege level in the status bar at the bottom of the main ASDM application window to ensure that you have the required privileges to execute privileged-level CLI commands.

Before You Begin

- Commands entered via the ASDM CLI tool might function differently from those entered through a terminal connection to the ASA.
- Command errors—If an error occurs because you entered an incorrect command, the incorrect command is skipped and the remaining commands are processed. A message appears in the Response area to inform you whether or not any error occurred, as well as other related information.
- Interactive commands—Interactive commands are not supported in the CLI tool. To use these commands in ASDM, use the **noconfirm** keyword if available, as shown in the following command:
`crypto key generate rsa modulus 1024 noconfirm`
- Avoid conflicts with other administrators—Multiple administrative users can update the running configuration of the ASA. Before using the ASDM CLI tool to make configuration changes, check for other active administrative sessions. If more than one user is configuring the ASA at the same time, the most recent changes take effect.

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

Procedure

-
- Step 1** In the main ASDM application window, choose **Tools > Command Line Interface**.
The **Command Line Interface** dialog box appears.
- Step 2** Choose the type of command (single line or multiple line) that you want, and then choose the command from the drop-down list, or type it in the field provided.
- Step 3** Click **Send** to execute the command.
- Step 4** To enter a new command, click **Clear Response**, and then choose (or type) another command to execute.
- Step 5** Check the **Enable context-sensitive help (?)** check box to provide context-sensitive help for this feature. Uncheck this check box to disable the context-sensitive help.
- Step 6** After you have closed the Command Line Interface dialog box, if you changed the configuration, click **Refresh** to view the changes in ASDM.
-

Show Commands Ignored by ASDM on the Device

This feature lets you show the list of commands that ASDM does not support. Typically, ASDM ignores them. ASDM does not change or remove these commands from your running configuration. See [Unsupported Commands, page 3-31](#) for more information.

Procedure

-
- Step 1** In the main ASDM application window, choose **Tools > Show Commands Ignored by ASDM on Device**.
- Step 2** Click **OK** when you are done.
-

Apply 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 that they can reconnect using the new policy.

To disconnect connections, enter one of the following commands:

- **clear local-host** *[ip_address]* **[all]**

This command reinitializes per-client run-time states such as connection limits and embryonic limits. As a result, this command removes any connection that uses those limits. See the **show local-host all** command to view all current connections per host.

With no arguments, this command clears all affected through-the-box connections. To also clear to-the-box connections (including your current management session), use the **all** keyword. To clear connections to and from a particular IP address, use the *ip_address* argument.

- **clear conn** [**all**] [**protocol** {**tcp** | **udp**}] [**address** *src_ip*[-*src_ip*] [**netmask** *mask*]] [**port** *src_port*[-*src_port*]] [**address** *dest_ip*[-*dest_ip*] [**netmask** *mask*]] [**port** *dest_port*[-*dest_port*]]

This command terminates connections in any state. See the **show conn** command to view all current connections.

With no arguments, this command clears all through-the-box connections. To also clear to-the-box connections (including your current management session), use the **all** keyword. To clear specific connections based on the source IP address, destination IP address, port, and/or protocol, you can specify the desired options.



ASDM Graphical User Interface

This chapter describes how to use the ASDM user interface.

- [About the ASDM User Interface, page 3-1](#)
- [Navigate the ASDM User Interface, page 3-3](#)
- [Menus, page 3-4](#)
- [Toolbar, page 3-9](#)
- [ASDM Assistant, page 3-10](#)
- [Status Bar, page 3-10](#)
- [Device List, page 3-11](#)
- [Common Buttons, page 3-11](#)
- [Keyboard Shortcuts, page 3-12](#)
- [Find Function in ASDM Panes, page 3-13](#)
- [Find Function in the ACL Manager Pane, page 3-14](#)
- [Enable Extended Screen Reader Support, page 3-15](#)
- [Organizational Folder, page 3-15](#)
- [Home Pane \(Single Mode and Context\), page 3-16](#)
- [Home Pane \(System\), page 3-28](#)
- [Define ASDM Preferences, page 3-29](#)
- [Search with the ASDM Assistant, page 3-31](#)
- [Enable History Metrics, page 3-31](#)
- [Unsupported Commands, page 3-31](#)

About the ASDM User Interface


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

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

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

The following figure shows the elements of the ASDM user interface.

Figure 3-1 ASDM User Interface



Legend

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

**Note**

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

Navigate the ASDM User Interface

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

- **Device Setup**
- **Firewall**
- **Botnet Traffic Filter**
- **Remote Access VPN**
- **Site to Site VPN**
- **Device Management**

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

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

-
- Step 1** Choose the drop-down list below the last function button to display a context menu.
- Step 2** Choose one of the following options:
- Click **Show More Buttons** to show more buttons.
 - Click **Show Fewer Buttons** to show fewer buttons.

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

Step 3 Click **OK** to save your settings and close this dialog box.

Menus

You can access ASDM menus using the mouse or keyboard. See [Keyboard Shortcuts, page 3-12](#) for information about accessing the menu bar from the keyboard.

ASDM has the following menus:

- [File Menu, page 3-4](#)
- [View Menu, page 3-5](#)
- [Tools Menu, page 3-6](#)
- [Wizards Menu, page 3-7](#)
- [Window Menu, page 3-8](#)
- [Help Menu, page 3-8](#)

File Menu

The **File** menu lets you manage ASA configurations.

File Menu Item	Description
Refresh ASDM with the Running Configuration on the Device	Loads a copy of the running configuration into ASDM.
Reset Device to the Factory Default Configuration	Restores the configuration to the factory default.
Show Running Configuration in New Window	Displays the current running configuration in a new window.
Save Running Configuration to Flash	Writes a copy of the running configuration to flash memory.
Save Running Configuration to TFTP Server	Stores a copy of the current running configuration file on a TFTP server.
Save Running Configuration to Standby Unit	Sends a copy of the running configuration file on the primary unit to the running configuration of a failover standby unit.
Save Internal Log Buffer to Flash	Saves the internal log buffer to flash memory.

File Menu Item	Description
Print	Prints the current page. We recommend landscape page orientation when you print rules. When you use Internet Explorer, permission to print was already granted when you originally accepted the signed applet.
Clear ASDM Cache	Removes local ASDM images. ASDM downloads images locally when you connect to ASDM.
Clear ASDM Password Cache	Removes the password cache if you have defined a new password and still have a existing password that is different than the new password.
Clear Internal Log Buffer	Empties the syslog message buffer.
Exit	Closes ASDM.

View Menu

The **View** menu lets you display various parts of the ASDM user interface. Certain items are dependent on the current view. You cannot select items that cannot be displayed in the current view.

View Menu Item	Description
Home	Displays the Home view.
Configuration	Displays the Configuration view.
Monitoring	Displays the Monitoring view.
Device List	Displays a list of devices in a dockable pane. See Device List, page 3-11 for more information.
Navigation	Shows and hides the display of the Navigation pane in the Configuration and Monitoring views.
ASDM Assistant	Searches and finds useful ASDM procedural help about certain tasks. See ASDM Assistant, page 3-10 for more information.
Latest ASDM Syslog Messages	Shows and hides the display of the Latest ASDM Syslog Messages pane in the Home view. This pane is only available in the Home view. If you do not have sufficient memory to upgrade to the most current release, syslog message %ASA-1-211004 is generated, indicating what the installed memory is and what the required memory is. This message reappears every 24 hours until the memory is upgraded.
Addresses	Shows and hides the display of the Addresses pane. The Addresses pane is only available for the Access Rules , NAT Rules , Service Policy Rules , AAA Rules , and Filter Rules panes in the Configuration view.
Services	Shows and hides the display of the Services pane. The Services pane is only available for the Access Rules , NAT Rules , Service Policy Rules , AAA Rules , and Filter Rules panes in the Configuration view.

View Menu Item	Description
Time Ranges	Shows and hides the display of the Time Ranges pane. The Time Ranges pane is only available for the Access Rules , Service Policy Rules , AAA Rules , and Filter Rules panes in the Configuration view.
Select Next Pane	Highlights the next pane shown in a multi-pane display, for example, going from the Service Policies Rules pane to the Address pane beside it.
Select Previous Pane	Highlights the previous pane shown in multi-pane displays.
Back	Returns to the previous pane.
Forward	Goes to the next pane previously visited.
Find in ASDM	Locates an item for which you are searching, such as a feature or the ASDM Assistant .
Reset Layout	Returns the layout to the default configuration.
Office Look and Feel	Changes the screen fonts and colors to the Microsoft Office settings.

Tools Menu

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

Tools Menu Item	Description
Command Line Interface	Sends commands to the ASA and view the results.
Show Commands Ignored by ASDM on Device	Displays unsupported commands that have been ignored by ASDM.
Packet Tracer	Traces a packet from a specified source address and interface to a destination. You can specify the protocol and port of any type of data and view the lifespan of a packet, with detailed information about actions taken on it. See the firewall configuration guide for more information.
Ping	Verifies the configuration and operation of the ASA and surrounding communications links, as well as performs basic testing of other network devices. See the firewall configuration guide for more information.
Traceroute	Determines the route that packets will take to their destination. See the firewall configuration guide for more information.
File Management	Views, moves, copies, and deletes files stored in flash memory. You can also create a directory in flash memory. You can also transfer files between various file systems, including TFTP, flash memory, and your local PC.
Check for ASA/ASDM Updates	Upgrades ASA software and ASDM software through a wizard.
Upgrade Software from Local Computer	Uploads an ASA image, ASDM image, or another image on your PC to flash memory.

Tools Menu Item	Description
Downgrade Software	Loads an older ASA image than the one you are currently running.
Backup Configurations	Backs up the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles.
Restore Configurations	Restores the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles.
System Reload	Restarts the ASDM and reload the saved configuration into memory.
Administrator's Alert to Clientless SSL VPN Users	Enables an administrator to send an alert message to clientless SSL VPN users. See the VPN configuration guide for more information.
Migrate Network Object Group Members	If you migrate to 8.3 or later, the ASA creates named network objects to replace inline IP addresses in some features. In addition to named objects, ASDM automatically creates non-named objects for any IP addresses used in the configuration. These auto-created objects are identified by the <i>IP address</i> only, do not have a name, and are not present as named objects in the platform configuration. When the ASA creates named objects as part of the migration, the matching non-named ASDM-only objects are replaced with the named objects. The only exception are non-named objects in a network object group. When the ASA creates named objects for IP addresses that are inside a network object group, ASDM retains the non-named objects as well, creating duplicate objects in ASDM. Choose Tools > Migrate Network Object Group Members to merge these object. See <i>Cisco ASA 5500 Migration to Version 8.3 and Later</i> for more information.
Preferences	Changes the behavior of specified ASDM functions between sessions. See Define ASDM Preferences, page 3-29 for more information.
ASDM Java Console	Shows the Java console.

Wizards Menu

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

Wizards Menu Item	Description
Startup Wizard	Guides you, step-by-step, through the initial configuration of the ASA.
VPN Wizards	There are separate wizards for a variety of VPN configurations. See the VPN configuration guide for more information.
High Availability and Scalability Wizard	Allows you to configure failover: VPN cluster load balancing, or ASA clustering on the ASA.

Wizards Menu Item	Description
Unified Communication Wizard	Enables you to configure unified communication features, such as an IP phone, on the ASA. See the firewall configuration guide for more information.
ASDM Identity Certificate Wizard	When using Java 7 update 51 and later, the ASDM Launcher requires a trusted certificate. An easy approach to fulfill the certificate requirements is to install a self-signed identity certificate. You can use Java Web Start to launch ASDM until you install a certificate using this wizard. See http://www.cisco.com/go/asdm-certificate for more information.
Packet Capture Wizard	Allows you to configure packet capture on the ASA. The wizard runs one packet capture on each ingress and egress interface. After you run the capture, you can save it on your computer, and then examine and analyze the capture with a packet analyzer.

Window Menu

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

Help Menu

The **Help** menu provides links to online help, as well as information about ASDM and the ASA.

Help Menu Items	Description
Help Topics	Opens a new browser window to show the ASDM online help. If you are managing an ASA FirePOWER module in ASDM, this item is labeled ASDM Help Topics .
ASA FirePOWER Help Topics	Opens a new browser window to show online help for the ASA FirePOWER module. This item is available only if you have installed the module and are managing it in ASDM.
Help for Current Screen	Opens context-sensitive help about the screen you are viewing. Alternatively, you can also click the ? Help button in the tool bar.
Release Notes	Opens the most current version of the <i>ASDM release notes</i> on Cisco.com. The release notes contain the most current information about ASDM software and hardware requirements, and the most current information about changes in the software.
Cisco ASA Series Documentation	Opens a document on Cisco.com that includes links to all of the available product documentation.
ASDM Assistant	Opens the ASDM Assistant , which lets you search downloadable content from Cisco.com, with details about performing certain tasks.

Help Menu Items	Description
About Cisco Adaptive Security Appliance (ASA)	Displays information about the ASA, including the software version, hardware set, configuration file loaded at startup, and software image loaded at startup. This information is helpful in troubleshooting.
About Cisco ASDM	Displays information about ASDM such as the software version, hostname, privilege level, operating system, device type, and Java version.

Toolbar

The **Toolbar** below the menus provides access to the Home view, Configuration view, and Monitoring view. It also lets you choose between the system and security contexts in multiple context mode, and provides navigation and other commonly used features.

Toolbar Button	Description
Home	Displays the Home pane, which lets you view important information about your ASA such as the status of your interfaces, the version you are running, licensing information, and performance. See Home Pane (Single Mode and Context), page 3-16 for more information. In multiple mode, the system does not have a Home pane.
Configuration	Configures the ASA. Click a function button in the left Navigation pane to configure that function.
Monitoring	Monitors the ASA. Click a function button in the left Navigation pane to configure that function.
Save Save ASA Changes	Saves the running configuration to the startup configuration for write-accessible contexts only. The button is labeled Save ASA Changes if you have an ASA FirePOWER module installed on the device and you are configuring it through ASDM.
Refresh	Refreshes ASDM with the current running configuration, except for graphs in any of the Monitoring panes.
Back	Returns to the last pane of ASDM that you visited.
Forward	Goes forward to the last pane of ASDM that you visited.
Help	Shows context-sensitive help for the screen that is currently open.
Search	Searches for a feature in ASDM. The Search function looks through the titles of each pane and presents you with a list of matches, and gives you a hyperlink directly to that pane. Click Back or Forward to switch quickly between two different panes that you found. See ASDM Assistant, page 3-10 for more information.

ASDM Assistant

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

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

To use the ASDM Assistant, perform the following steps:

-
- Step 1** Choose **View > ASDM Assistant**.
The **ASDM Assistant** pane appears.
- Step 2** Enter the information that you want to find in the **Search** field, then click **Go**.
The requested information appears in the **Search Results** pane.
- Step 3** Click any links that appear in the **Search Results and Features** areas to obtain more details.
-

Status Bar

The **Status Bar** appears at the bottom of the ASDM window. The following table lists the areas shown from left to right.

Area	Description
Status	The status of the configuration (for example, “Device configuration loaded successfully.”)
Failover	The status of the failover unit, either active or standby.
User Name	The username of the ASDM user. If you logged in without a username, the username is “admin.”
User Privilege	The privilege of the ASDM user.
Commands Ignored by ASDM	Click the icon to show a list of commands from your configuration that ASDM did not process. These commands will not be removed from the configuration.
Connection to Device	The ASDM connection status to the ASA. See Connection to Device, page 3-11 for more information.
Syslog Connection	The syslog connection is up, and the ASA is being monitored.
SSL Secure	The connection to ASDM is secure because it uses SSL.
Time	The time that is set on the ASA.

Connection to Device

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

Device List

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

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

-
- Step 1** Click **Add** to add another device to the list.
The **Add Device** dialog box appears.
 - Step 2** Enter the device name or IP address of the device, then click **OK**.
 - Step 3** Click **Delete** to remove a selected device from the list.
 - Step 4** Click **Connect** to connect to another device.
The **Enter Network Password** dialog box appears.
 - Step 5** Enter your username and password in the applicable fields, then click **Login**.
-

Common Buttons

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

Button	Description
Apply	Sends changes made in ASDM to the ASA and applies them to the running configuration.
Save	Writes a copy of the running configuration to flash memory.
Reset	Discards changes and reverts to the information displayed before changes were made or the last time that you clicked Refresh or Apply. After you click Reset , click Refresh to make sure that information from the current running configuration appears.
Restore Default	Clears the selected settings and returns to the default settings.
Cancel	Discards changes and returns to the previous pane.
Enable	Displays read-only statistics for a feature.

Button	Description
Close	Closes an open dialog box.
Clear	Remove information from a field, or remove a check from a check box.
Back	Returns to the previous pane.
Forward	Goes to the next pane.
Help	Displays help for the selected pane or dialog box.

Keyboard Shortcuts

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

The following table lists the keyboard shortcuts you can use to move across the three main areas of the ASDM user interface.

Table 3-1 *Keyboard Shortcuts Within the Main Window*

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

The following table lists the keyboard shortcut you can use to navigate within a pane.

Table 3-2 *Keyboard Shortcuts Within a Pane*

To move the focus to the	Press
Next field	Tab
Previous field	Shift+Tab

Table 3-2 Keyboard Shortcuts Within a Pane (continued)

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

The following table lists the keyboard shortcuts you can use with the Log Viewers.

Table 3-3 Keyboard Shortcuts for the Log Viewer

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

The following table lists the keyboard shortcuts you can use to access menu items.

Table 3-4 Keyboard Shortcuts to Access Menu Items

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

Find Function in ASDM Panes

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

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

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

Boston-LA, Boston-Lisbon, Boston-London

Entering Bo?ton might return the following matches:

Boston, Bolton

Find Function in the ACL Manager Pane

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

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

-
- Step 1** Click **Find** in the **ACL Manager** pane.
- Step 2** Choose one of the following options in the **Filter** field from the drop-down list:
- **Source**—The search includes a source IP address of a the network object group, interface IP, or any address from which traffic is permitted or denied. You specify this address in [Step 4](#).
 - **Destination**—The search includes a destination IP address (host or network) that is permitted or denied to send traffic to the IP addresses listed in the **Source** section. You specify this address in [Step 4](#).
 - **Source or Destination**—The search includes either a source or a destination address that you specify in [Step 4](#).
 - **Service**—The search includes a service group or predefined service policy that you specify in [Step 4](#).
 - **Query**—When you choose **Query** from the drop-down list, click **Query** to specify a detailed search by all four of the preceding options: **Source**, **Destination**, **Source or Destination**, and **Service**.
- Step 3** In the second field, choose one of the following options from the drop-down list:
- **is**—Specifies an exact match of the detail that you enter in [Step 4](#).
 - **contains**—Specifies to search for ACLs or ACEs that contain, but are not limited to, the detail that you enter in [Step 4](#).
- Step 4** In the third field, enter specific criteria about ACLs or ACEs that you would like to find, or click **Browse** to search for key elements in your ACL/ACE configuration.
- Step 5** Click **Filter** to perform the search.
- The ASDM find function returns a list of ACLs and ACEs that contain your specified criteria.
- Step 6** Click **Clear** to clear the list of found ACLs and ACEs.
- Step 7** Click the red **x** to close the find function dialog box.
-

Enable Extended Screen Reader Support

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

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

-
Step 1	Choose Tools > Preferences . The Preferences dialog box appears.
Step 2	Check the Enable screen reader support check box on the General tab.
Step 3	Click OK .
Step 4	Restart ASDM to activate screen reader support.
-

Organizational Folder

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

Home Pane (Single Mode and Context)

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


If you have hardware or software modules installed on the device, such as IPS, CX, or ASA FirePOWER modules, there are separate tabs for those modules.

Device Dashboard Tab

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

The following figure shows the elements of the **Device Dashboard** tab.

Figure 3-2 Device Dashboard Tab



Legend

GUI Element	Description
1	Device Information Pane, page 3-17
2	Interface Status Pane, page 3-18
3	VPN Sessions Pane, page 3-18
4	Traffic Status Pane, page 3-19

GUI Element	Description
5	System Resources Status Pane, page 3-18
6	Traffic Status Pane, page 3-19
—	Device List, page 3-11
—	Latest ASDM Syslog Messages Pane, page 3-19

Device Information Pane

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

General Tab

This tab shows basic information about the ASA:

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

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

See the hardware guide for your particular device for more information about specific hardware statistics.



Note


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

License Tab

This tab shows a subset of licensed features. Click **More Licenses** to view detailed license information, or to enter a new activation key; the **Configuration > Device Management > Licensing > Activation Key pane** appears.

Cluster Tab

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



Virtual Resources Tab (ASAv)

This tab shows the virtual resources used by the ASAv, including the number of vCPUs, RAM, and whether the ASAv is over- or under-provisioned.

Interface Status Pane

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

VPN Sessions Pane

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

Failover Status Pane

This pane shows the failover status.

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

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

System Resources Status Pane

This pane shows CPU and memory usage statistics.

Traffic Status Pane

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


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

Latest ASDM Syslog Messages Pane

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

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

Figure 3-3 Latest ASDM Syslog Messages Pane



Legend

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

- Right-click an event and choose **Clear Content** to clear the current messages.
- Right-click an event and click **Save Content** to save the current messages to a file on your PC.
- Right-click an event and choose **Copy** to copy the current content.

- Right-click an event and choose **Color Settings** to change the background and foreground colors of syslog messages according to their severity.

Firewall Dashboard Tab

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

Figure 3-4 shows some of the elements of the **Firewall Dashboard** tab.

Figure 3-4 Firewall Dashboard Tab



Legend

GUI Element	Description
1	Traffic Overview Pane, page 3-21
2	Top 10 Access Rules Pane, page 3-21
3	Top Usage Status Pane, page 3-21
(not shown)	Top Ten Protected Servers Under SYN Attack Pane, page 3-22
(not shown)	Top 200 Hosts Pane, page 3-22
(not shown)	Top Botnet Traffic Filter Hits Pane, page 3-22

Traffic Overview Pane

Enabled by default. If you disable basic threat detection (see the firewall configuration guide), then this area includes an **Enable** button that lets you enable basic threat detection. The runtime statistics include the following information, which is *display-only*:

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

Top 10 Access Rules Pane

Enabled by default. If you disable threat detection statistics for access rules (see the firewall configuration guide), then this area includes an **Enable** button that lets you enable statistics for access rules.

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

Top Usage Status Pane

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

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

The first three tabs—**Top 10 Services**, **Top 10 Sources**, and **Top 10 Destinations**—provide statistics for threat detection services. Each tab includes an **Enable** button that let you enable each threat detection service. You can enable them according to the firewall configuration guide.

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

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

The **Top 10 Users** tab displays data only when you have configured the Identity Firewall service in the ASA, which includes configuring these additional components—Microsoft Active Directory and Cisco Active Directory (AD) Agent.

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

**Caution**

Enabling statistics can affect the ASA performance, depending on the type of statistics enabled. Enabling statistics for hosts affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Enabling statistics for ports, however, has a modest effect.

Top Ten Protected Servers Under SYN Attack Pane

Disabled by default. This area includes an **Enable** button that lets you enable the feature, or you can enable it according to the firewall configuration guide. Statistics for the top ten protected servers under attack are displayed.

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

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

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

Top 200 Hosts Pane

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


Top Botnet Traffic Filter Hits Pane

Disabled by default. This area includes links to configure the Botnet Traffic Filter. Reports of the top ten botnet sites, ports, and infected hosts provide a snapshot of the data, and may not match the top ten items since statistics started to be collected. If you right-click an IP address, you can invoke the whois tool to learn more about the botnet site.

See the firewall configuration guide for more information.

Cluster Dashboard Tab

When you enable ASA clustering and are connected to the master unit, the **Cluster Dashboard** tab shows a summary of cluster membership and resource utilization.



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



Note

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

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

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

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

– **Throughput** tab:

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

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

– **Load Balancing** tab:

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

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


– **Control Link Usage** tab:

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

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

Cluster Firewall Dashboard Tab

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




Intrusion Prevention Tab

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

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

- Step 1** Click the **Intrusion Prevention** tab.
- The **Connecting to IPS** dialog box appears.




- Step 2** Enter the IP address, port, username and password. The default IP address and port is 192.168.1.2:443. The default username and password is **cisco** and **cisco**.
- Step 3** Check the **Save IPS login** information on local host check box to save the login information on your local PC.
- Step 4** Click **Continue**.

See the firewall configuration guide for more information about intrusion prevention.

The following figure shows the elements of the **Health Dashboard** tab, located on the **Intrusion Prevention** tab.

Figure 3-5 Intrusion Prevention Tab (Health Dashboard)




Legend

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

ASA CX Status Tab

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



ASA FirePOWER Tabs

The **ASA FirePOWER Status** tab lets you view information about the module. This includes module information, such as the model, serial number, and software version, and module status, such as the application name and status, data plane status, and overall status. If the module is registered to a FireSIGHT Management Center, you can click the link to open the application and do further analysis and module configuration.

This tab appears only if you have an ASA FirePOWER module installed in the device.

If you are managing the ASA FirePOWER module with ASDM rather than FireSIGHT Management Center, there are additional tabs:

- **ASA FirePOWER Dashboard**—The dashboard provides summary information about the software running on the module, product updates, licensing, system load, disk usage, system time, and interface status.


- ASA FirePOWER Reporting**—The reporting page provides Top 10 dashboards for a wide variety of module statistics, such as web categories, users, sources, and destinations for the traffic passing through the module.

Home Pane (System)

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

The following figure shows the elements of the System **Home** pane.

Figure 3-6 System Home Pane



Legend

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

GUI Element	Description
4	CPU Status pane.
5	Memory Status pane.

Define ASDM Preferences

You can define the behavior of certain ASDM settings.

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

Step 1 Choose **Tools > Preferences**.

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

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

Step 3 On the **General** tab, specify the following:

- Check the **Warn that configuration in ASDM is out of sync with the configuration in ASA** check box to be notified when the startup configuration and the running configuration are no longer in sync with each other.
- Check the **Show configuration restriction message to read-only user** check box to display the following message to a read-only user at startup. This option is checked by default.
"You are not allowed to modify the ASA configuration, because you do not have sufficient privileges."
- Check the **Show configuration restriction message on a slave unit in an ASA cluster** check box to display a configuration restriction message to a user connected to a slave unit.
- Check the **Confirm before exiting ASDM** check box to display a prompt when you try to close ASDM to confirm that you want to exit. This option is checked by default.
- Check the **Enable screen reader support (requires ASDM restart)** check box to enable screen readers to work. You must restart ASDM to enable this option.
- Check the **Warn of insufficient ASA memory when ASDM loads** check box to receive notification when the minimum amount of ASA memory is insufficient to run complete functionality in the ASDM application. ASDM displays the memory warning in a text banner message at bootup, displays a message in the title bar text in ASDM, and sends a syslog alert once every 24 hours.
- In the **Communications** area:
 - Check the **Preview commands before sending them to the device** check box to view CLI commands generated by ASDM.
 - Check the **Enable cumulative (batch) CLI delivery** check box to send multiple commands in a single group to the ASA.
 - Enter the minimum amount of time in seconds for a configuration to send a timeout message. The default is 60 seconds.
- In the **Logging** area:
 - Check the **Enable logging to the ASDM Java console** check box to configure Java logging.

- Set the severity level by choosing a **Logging Level** from the drop-down list.
- In the **Packet Capture Wizard** area, to display captured packets, enter the name of the **Network Sniffer Application** or click **Browse** to find it in the file system.

Step 4 On the **Rules Table** tab, specify the following:

- Display settings let you change the way rules appear in the Rules table.
 - Check the **Auto-expand network and service object groups with specified prefix** check box to display the network and service object groups automatically expanded based on the Auto-Expand Prefix setting.
 - Enter the prefix of the network and service object groups to expand automatically when displayed in the **Auto-Expand Prefix** field.
 - Check the **Show members of network and service object groups** check box to display members of network and service object groups and the group name in the Rules table. If the check box is not checked, only the group name is displayed.
 - Enter the number of network and service object groups to display in the **Limit Members To** field. When the object group members are displayed, then only the first *n* members are displayed.
 - Check the **Show all actions for service policy rules** check box to display all actions in the Rules table. When unchecked, a summary appears.
- Deployment settings let you configure the behavior of the ASA when deploying changes to the Rules table.
 - Check the **Issue “clear xlate” command when deploying access lists** check box to clear the NAT table when deploying new access lists. This setting ensures the access lists that are configured on the ASA are applied to all translated addresses.
- Access Rule Hit Count Settings let you configure the frequency for which the hit counts are updated in the Access Rules table. Hit counts are applicable for explicit rules only. No hit count will be displayed for implicit rules in the Access Rules table.
 - Check the **Update access rule hit counts automatically** check box to have the hit counts automatically updated in the Access Rules table.
 - Specify the frequency in seconds in which the hit count column is updated in the Access Rules table. Valid values are 10 - 86400 seconds.

Step 5 On the **Syslog** tab, specify the following:

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

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

**Note**

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

Search with the ASDM Assistant

The ASDM Assistant tool lets you search and view useful ASDM procedural help about certain tasks. Choose **View > ASDM Assistant > How Do I?** to access information, or enter a search request from the **Look For** field in the menu bar. Choose **How Do I?** From the **Find** drop-down list to begin the search. To view the ASDM Assistant, perform the following steps:

-
- Step 1** Choose **View > ASDM Assistant**.
The **ASDM Assistant** pane appears.
 - Step 2** Enter the information that you want to find in the **Search** field, and click **Go**.
The requested information appears in the **Search Results** pane.
 - Step 3** Click any links that appear in the **Search Results and Features** sections to obtain more details.
-

Enable History Metrics

The **Configuration > Device Management > Advanced > History Metrics** pane lets you configure the ASA to keep a history of various statistics, which ASDM can display on any graph/table. If you do not enable history metrics, you can only monitor statistics in real time. Enabling history metrics lets you view statistics graphs from the last 10 minutes, 60 minutes, 12 hours, and 5 days.

To configure history metrics, perform the following steps:

-
- Step 1** Choose **Configuration > Device Management > Advanced > History Metrics**.
The **History Metrics** pane appears.
 - Step 2** Check the **ASDM History Metrics** check box to enable history metrics, then click **Apply**.
-

Unsupported Commands

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

Ignored and View-Only Commands

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

Table 3-5 *List of Unsupported Commands*

Unsupported Commands	ASDM Behavior
capture	Ignored.
coredump	Ignored. This can be configured only using the CLI.
crypto engine large-mod-accel	Ignored.
dhcp-server (tunnel-group name general-attributes)	ASDM only allows one setting for all DHCP servers.
eject	Unsupported.
established	Ignored.
failover timeout	Ignored.
fips	Ignored.
nat-assigned-to-public-ip	Ignored.
pager	Ignored.
pim accept-register route-map	Ignored. You can configure only the list option using ASDM.
service-policy global	Ignored if it uses a match access-list class. For example: <pre>access-list myacl extended permit ip any any class-map mycm match access-list myacl policy-map mypm class mycm inspect ftp service-policy mypm global</pre>
set metric	Ignored.
sysopt nodnsalias	Ignored.
sysopt uauth allow-http-cache	Ignored.
terminal	Ignored.
threat-detection rate	Ignored.

Effects of Unsupported Commands

If ASDM loads an existing running configuration and finds other unsupported commands, ASDM operation is unaffected. Choose **Tools > Show Commands Ignored by ASDM on Device** to view the unsupported commands.

Discontinuous Subnet Masks Not Supported

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

```
ip address inside 192.168.2.1 255.255.0.255
```

Interactive User Commands Not Supported by the ASDM CLI Tool

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

For example:

1. Choose **Tools > Command Line Interface**.
2. Enter the **crypto key generate rsa** command.
ASDM generates the default 1024-bit RSA key.
3. Enter the **crypto key generate rsa** command again.

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

```
Do you really want to replace them? [yes/no]:WARNING: You already have RSA
ke00000000000000$A key
Input line must be less than 16 characters in length.

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

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

Workaround:

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

```
crypto key generate rsa noconfirm
```




Product Authorization Key Licenses

A license specifies the options that are enabled on a given Cisco ASA. This document describes product authorization key (PAK) licenses for all physical ASAs. For the ASAv, see [Chapter 4, “Smart Software Licensing for the ASAv.”](#)

- [Supported Feature Licenses Per Model, page 3-1](#)
- [About PAK Licenses, page 3-19](#)
- [Guidelines for PAK Licenses, page 3-27](#)
- [Configure PAK Licenses, page 3-28](#)
- [Configure a Shared License \(AnyConnect 3 and Earlier\), page 3-30](#)
- [Monitoring PAK Licenses, page 3-35](#)
- [History for PAK Licenses, page 3-36](#)

Supported Feature Licenses Per Model

This section describes the licenses available for each model as well as important notes about licenses.

- [Licenses Per Model, page 3-1](#)
- [License Notes, page 3-15](#)

Licenses Per Model

This section lists the feature licenses available for each model:

- [ASA 5506-X, ASA 5506W-X, and ASA 5506H-X, page 3-3](#)
- [ASA 5508-X, page 3-3](#)
- [ASA 5512-X, page 3-4](#)
- [ASA 5515-X, page 3-5](#)
- [ASA 5516-X, page 3-6](#)
- [ASA 5525-X, page 3-8](#)
- [ASA 5545-X, page 3-8](#)
- [ASA 5555-X, page 3-9](#)
- [ASA 5585-X with SSP-10, page 3-10](#)

- [ASA 5585-X with SSP-20, page 3-11](#)
- [ASA 5585-X with SSP-40 and -60, page 3-12](#)
- [ASA Services Module, page 3-13](#)

Items that are in *italics* are separate, optional licenses that can replace the Base (or Security Plus, and so on) 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 [No Payload Encryption Models, page 3-26](#) for a list of unsupported features.

For detailed information about licenses, see [License Notes, page 3-15](#).

ASA 5506-X, ASA 5506W-X, and ASA 5506H-X**Table 3-1 ASA 5506-X, ASA 5506W-X, and ASA 5506H-X License Features**

Licenses	Base License		Security Plus License	
Firewall Licenses				
Botnet Traffic Filter	No support		Enabled	
Firewall Conns, Concurrent	20,000		50,000	
GTP/GPRS	No Support		No Support	
Total UC Proxy Sessions	160		160	
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.				
AnyConnect Premium Peers (maximum)	50		50	
	Shared licenses: No Support		Shared licenses: No Support	
Adv. Endpoint Assessment	Enabled		Enabled	
AnyConnect for Cisco VPN Phone	Enabled		Enabled	
AnyConnect Essentials	No support		No support	
AnyConnect for Mobile	Enabled		Enabled	
Other VPN Licenses				
Total VPN Peers, combined all types	50		50	
Other VPN Peers	10		50	
VPN Load Balancing	No support		No support	
General Licenses				
Encryption	Base (DES)	Opt. lic.: Strong (3DES/AES)	Base (DES)	Opt. lic.: Strong (3DES/AES)
Failover	No support		Active/Standby	
Interfaces of all types, Max.	536		636	
Security Contexts	No support		No support	
Clustering	No Support		No Support	
VLANs, Maximum	5		30	

ASA 5508-X**Table 3-2 ASA 5508-X License Features**

Licenses	Base License
Firewall Licenses	
Botnet Traffic Filter	Enabled
Firewall Conns, Concurrent	100,000
GTP/GPRS	No Support
Total UC Proxy Sessions	320

Table 3-2 ASA 5508-X License Features (continued)

Licenses	Base License		
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.			
AnyConnect Premium Peers (maximum)	100		
	Shared licenses: No Support		
Adv. Endpoint Assessment	Enabled		
AnyConnect for Cisco VPN Phone	Enabled		
AnyConnect Essentials	No support		
AnyConnect for Mobile	Enabled		
Other VPN Licenses			
Total VPN Peers, combined all types	100		
Other VPN Peers	100		
VPN Load Balancing	Enabled		
General Licenses			
Encryption	Base (DES)	Opt. lic.: Strong (3DES/AES)	
Failover	Active/Standby or Active/Active		
Interfaces of all types, Max.	716		
Security Contexts	2	Optional licenses:	5
Clustering	No Support		
VLANs, Maximum	50		

ASA 5512-X**Table 3-3 ASA 5512-X License Features**

Licenses	Base License						Security Plus License					
Firewall Licenses												
Botnet Traffic Filter	Disabled		Optional Time-based license: Available				Disabled		Optional Time-based license: Available			
Firewall Conns, Concurrent	100,000						250,000					
GTP/GPRS	No support						Disabled		Optional license: Available			
Total UC Proxy Sessions	2	Optional licenses:					2	Optional licenses:				
		24	50	100	250	500		24	50	100	250	500

VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.

Table 3-3 ASA 5512-X License Features (continued)

Licenses	Base License		Security Plus License		
AnyConnect Premium Peers (Maximum)	250		250		
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.		Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.		
Adv. Endpoint Assessment	Enabled		Enabled		
AnyConnect for Cisco VPN Phone	Enabled		Enabled		
AnyConnect Essentials	Disabled		Disabled		
AnyConnect for Mobile	Enabled		Enabled		
Other VPN Licenses					
Total VPN Peers, combined all types	250		250		
Other VPN Peers	250		250		
VPN Load Balancing	No support		Enabled		
General Licenses					
Encryption	Base (DES)	Opt. lic.: Strong (3DES/AES)	Base (DES)	Opt. lic.: Strong (3DES/AES)	
Failover	No support		Active/Standby or Active/Active		
Interfaces of all types, Max.	716		916		
Security Contexts	No support		2	Optional licenses:	5
Clustering	No Support		2		
IPS Module	Disabled	Optional license: Available	Disabled	Optional license: Available	
VLANs, Maximum	50		100		

ASA 5515-X**Table 3-4 ASA 5515-X License Features**

Licenses	Base License							
Firewall Licenses								
Botnet Traffic Filter	Disabled		Optional Time-based license: Available					
Firewall Conns, Concurrent	250,000							
GTP/GPRS	Disabled		Optional license: Available					
Total UC Proxy Sessions	2	Optional licenses:		24	50	100	250	500
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.								
AnyConnect Premium Peers (maximum)	250							
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.							
Adv. Endpoint Assessment	Enabled							

Table 3-4 ASA 5515-X License Features (continued)

Licenses	Base License		
AnyConnect for Cisco VPN Phone	Enabled		
AnyConnect Essentials	Disabled		
AnyConnect for Mobile	Enabled		
Other VPN Licenses			
Total VPN Peers, combined all types	250		
Other VPN Peers	250		
VPN Load Balancing	Enabled		
General Licenses			
Encryption	Base (DES)	Optional license: Strong (3DES/AES)	
Failover	Active/Standby or Active/Active		
Interfaces of all types, Max.	916		
Security Contexts	2	Optional licenses:	5
Clustering	2		
IPS Module	Disabled	Optional license: Available	
VLANs, Maximum	100		

ASA 5516-X**Table 3-5 ASA 5516-X License Features**

Licenses	Base License	
Firewall Licenses		
Botnet Traffic Filter	Enabled	
Firewall Conns, Concurrent	250,000	
GTP/GPRS	Disabled	Optional license: Available
Total UC Proxy Sessions	1000	
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.		
AnyConnect Premium Peers (maximum)	300	
	Shared licenses: No Support	
Adv. Endpoint Assessment	Enabled	
AnyConnect for Cisco VPN Phone	Enabled	
AnyConnect Essentials	No Support	
AnyConnect for Mobile	Enabled	
Other VPN Licenses		

Table 3-5 ASA 5516-X License Features (continued)

Licenses	Base License		
Total VPN Peers, combined all types	300		
Other VPN Peers	300		
VPN Load Balancing	Enabled		
General Licenses			
Encryption	Base (DES)	Opt. lic.: Strong (3DES/AES)	
Failover	Active/Standby or Active/Active		
Interfaces of all types, Max.	1,116		
Security Contexts	2	Optional licenses:	5
Clustering	No Support		
VLANs, Maximum	150		

ASA 5525-X**Table 3-6 ASA 5525-X License Features**

Licenses	Base License								
Firewall Licenses									
Botnet Traffic Filter	Disabled		Optional Time-based license: Available						
Firewall Conns, Concurrent	500,000								
GTP/GPRS	Disabled		Optional license: Available						
Total UC Proxy Sessions	2	Optional licenses:	24	50	100	250	500	750	1000
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.									
AnyConnect Premium Peers (maximum)	750								
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.								
Adv. Endpoint Assessment	Enabled								
AnyConnect for Cisco VPN Phone	Enabled								
AnyConnect Essentials	Disabled								
AnyConnect for Mobile	Enabled								
Other VPN Licenses									
Total VPN Peers, combined all types	750								
Other VPN Peers	750								
VPN Load Balancing	Enabled								
General Licenses									
Encryption	Base (DES)		Optional license: Strong (3DES/AES)						
Failover	Active/Standby or Active/Active								
Interfaces of all types, Max.	1316								
Security Contexts	2	Optional licenses:	5	10	20				
Clustering	2								
IPS Module	Disabled		Optional license: Available						
VLANs, Maximum	200								

ASA 5545-X**Table 3-7 ASA 5545-X License Features**

Licenses	Base License	
Firewall Licenses		
Botnet Traffic Filter	Disabled	Optional Time-based license: Available
Firewall Conns, Concurrent	750,000	
GTP/GPRS	Disabled	Optional license: Available

Table 3-7 ASA 5545-X License Features (continued)

Licenses	Base License									
Total UC Proxy Sessions	2	Optional licenses:	24	50	100	250	500	750	1000	2000
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.										
AnyConnect Premium Peers (maximum)	2500									
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.									
Adv. Endpoint Assessment	Enabled									
AnyConnect for Cisco VPN Phone	Enabled									
AnyConnect Essentials	Disabled									
AnyConnect for Mobile	Enabled									
Other VPN Licenses										
Total VPN Peers, combined all types	2500									
Other VPN Peers	2500									
VPN Load Balancing	Enabled									
General Licenses										
Encryption	Base (DES)	Optional license: Strong (3DES/AES)								
Failover	Active/Standby or Active/Active									
Interfaces of all types, Max.	1716									
Security Contexts	2	Optional licenses:	5	10	20	50				
Clustering	2									
IPS Module	Disabled	Optional license: Available								
VLANs, Maximum	300									

ASA 5555-X**Table 3-8 ASA 5555-X License Features**

Licenses	Base License									
Firewall Licenses										
Botnet Traffic Filter	Disabled		Optional Time-based license: Available							
Firewall Conns, Concurrent	1,000,000									
GTP/GPRS	Disabled		Optional license: Available							
Total UC Proxy Sessions	2	Optional licenses:								
		24	50	100	250	500	750	1000	2000	3000
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.										

Table 3-8 ASA 5555-X License Features (continued)

Licenses	Base License							
AnyConnect Premium Peers (maximum)	5000							
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.							
Adv. Endpoint Assessment	Enabled							
AnyConnect for Cisco VPN Phone	Enabled							
AnyConnect Essentials	Disabled							
AnyConnect for Mobile	Enabled							
Other VPN Licenses								
Total VPN Peers, combined all types	5000							
Other VPN Peers	5000							
VPN Load Balancing	Enabled							
General Licenses								
Encryption	Base (DES)		Optional license: Strong (3DES/AES)					
Failover	Active/Standby or Active/Active							
Interfaces of all types, Max.	2516							
Security Contexts	2	Optional licenses:	5	10	20	50	100	
Clustering	2							
IPS Module	Disabled		Optional license: Available					
VLANs, Maximum	500							

ASA 5585-X with SSP-10

You can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-10 with an SSP-20 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired.

Table 3-9 ASA 5585-X with SSP-10 License Features

Licenses	Base and Security Plus Licenses									
Firewall Licenses										
Botnet Traffic Filter	Disabled	Optional Time-based license: Available								
Firewall Conns, Concurrent	1,000,000									
GTP/GPRS	Disabled	Optional license: Available								
Total UC Proxy Sessions	2	Optional licenses:								
	24	50	100	250	500	750	1000	2000	3000	

VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.

Table 3-9 ASA 5585-X with SSP-10 License Features (continued)

Licenses	Base and Security Plus Licenses							
AnyConnect Premium Peers (maximum)	5000							
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.							
Adv. Endpoint Assessment	Enabled							
AnyConnect for Cisco VPN Phone	Enabled							
AnyConnect Essentials	Disabled							
AnyConnect for Mobile	Enabled							
Other VPN Licenses								
Total VPN Peers, combined all types	5000							
Other VPN Peers	5000							
VPN Load Balancing	Enabled							
General Licenses								
10 GE I/O	Base License: Disabled; fiber ifcs run at 1 GE				Security Plus License: Enabled; fiber ifcs run at 10 GE			
Encryption	Base (DES)	Optional license: Strong (3DES/AES)						
Failover	Active/Standby or Active/Active							
Interfaces of all types, Max.	4612							
Security Contexts	2	Optional licenses:	5	10	20	50	100	
Clustering	Disabled		Optional license: Available for 16 units					
VLANs, Maximum	1024							

ASA 5585-X with SSP-20

You can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-20 with an SSP-40 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired.

Table 3-10 ASA 5585-X with SSP-20 License Features

Licenses	Base and Security Plus Licenses											
Firewall Licenses												
Botnet Traffic Filter	Disabled	Optional Time-based license: Available										
Firewall Conns, Concurrent	2,000,000											
GTP/GPRS	Disabled	Optional license: Available										
Total UC Proxy Sessions	2	Optional licenses:										
		24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹

VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.

Table 3-10 ASA 5585-X with SSP-20 License Features (continued)

Licenses	Base and Security Plus Licenses								
AnyConnect Premium Peers (maximum)	10,000								
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.								
Adv. Endpoint Assessment	Enabled								
AnyConnect for Cisco VPN Phone	Enabled								
AnyConnect Essentials	Disabled								
AnyConnect for Mobile	Enabled								
Other VPN Licenses									
Total VPN Peers, combined all types	10,000								
Other VPN Peers	10,000								
VPN Load Balancing	Enabled								
General Licenses									
10 GE I/O	Base License: Disabled; fiber ifcs run at 1 GE					Security Plus License: Enabled; fiber ifcs run at 10 GE			
Encryption	Base (DES)	Optional license: Strong (3DES/AES)							
Failover	Active/Standby or Active/Active								
Interfaces of all types, Max.	4612								
Security Contexts	2	Optional licenses:	5	10	20	50	100	250	
Clustering	Disabled		Optional license: Available for 16 units						
VLANs, Maximum	1024								

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

ASA 5585-X with SSP-40 and -60

You can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired.

Table 3-11 ASA 5585-X with SSP-40 and -60 License Features

Licenses	Base License											
Firewall Licenses												
Botnet Traffic Filter	Disabled		Optional Time-based license: Available									
Firewall Conns, Concurrent	5585-X with SSP-40: 4,000,000							5585-X with SSP-60: 10,000,000				
GTP/GPRS	Disabled		Optional license: Available									
Total UC Proxy Sessions	2	Optional licenses:										
		24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹

VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.

Table 3-11 ASA 5585-X with SSP-40 and -60 License Features (continued)

Licenses	Base License								
AnyConnect Premium Peers (maximum)	10,000								
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.								
Adv. Endpoint Assessment	Enabled								
AnyConnect for Cisco VPN Phone	Enabled								
AnyConnect Essentials	Disabled								
AnyConnect for Mobile	Enabled								
Other VPN Licenses									
Total VPN Peers, combined all types	10,000								
Other VPN Peers	10,000								
VPN Load Balancing	Enabled								
General Licenses									
10 GE I/O	Enabled; fiber ifcs run at 10 GE								
Encryption	Base (DES)	Optional license: Strong (3DES/AES)							
Failover	Active/Standby or Active/Active								
Interfaces of all types, Max.	4612								
Security Contexts	2	Optional licenses:	5	10	20	50	100	250	
Clustering	Disabled		Optional license: Available for 16 units						
VLANs, Maximum	1024								

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

ASA Services Module

Table 3-12 ASASM License Features

Licenses	Base License											
Firewall Licenses												
Botnet Traffic Filter	Disabled		Optional Time-based license: Available									
Firewall Conns, Concurrent	10,000,000											
GTP/GPRS	Disabled		Optional license: Available									
Total UC Proxy Sessions	2	Optional licenses:										
		24	50	100	250	500	750	1000	2000	3000	5000	10,000 ¹
VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.												
AnyConnect Premium Peers (maximum)	10,000											
	Optional Shared licenses: We recommend upgrading to the AnyConnect Plus or Apex licensing, which replaces shared server functionality.											

Table 3-12 ASASM License Features (continued)

Licenses	Base License						
Adv. Endpoint Assessment	Enabled						
AnyConnect for Cisco VPN Phone	Enabled						
AnyConnect Essentials	Disabled						
AnyConnect for Mobile	Enabled						
Other VPN Licenses							
Total VPN Peers, combined all types	10,000						
Other VPN Peers	10,000						
VPN Load Balancing	Enabled						
General Licenses							
Encryption	Base (DES)		Optional license: Strong (3DES/AES)				
Failover	Active/Standby or Active/Active						
Security Contexts	2	Optional licenses:					
		5	10	20	50	100	250
Clustering	No support						
VLANs, Maximum	1000						

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

License Notes

The following table includes additional information about licenses.

Table 3-13 *License Notes*

License	Notes
AnyConnect Essentials	Note This license is a legacy license. We recommend that you upgrade to the AnyConnect Plus or Apex license. The AnyConnect Essentials sessions include the following VPN types: • SSL VPN • IPsec remote access VPN using IKEv2 This license does not support browser-based (clientless) SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium license instead of the AnyConnect Essentials license. Note With the AnyConnect Essentials license, VPN users can use a web browser to log in, and download and start (WebLaunch) the AnyConnect client. The AnyConnect client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given ASA: AnyConnect Premium license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium licenses on different ASAs in the same network. By default, the ASA uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the webvpn, and then the no anyconnect-essentials command or in ASDM, using the Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials pane.
AnyConnect for Cisco VPN Phone	Note This license is a legacy license. We recommend that you upgrade to the AnyConnect Plus or Apex license, which includes this license. In conjunction with an AnyConnect Premium license, this license enables access from hardware IP phones that have built in AnyConnect compatibility.

Table 3-13 License Notes (continued)

License	Notes
AnyConnect for Mobile	Note This license is a legacy license. We recommend that you upgrade to the AnyConnect Plus or Apex license, which includes this license. This license provides access to the AnyConnect Client for touch-screen mobile devices running Windows Mobile 5.0, 6.0, and 6.1. We recommend using this license if you want to support mobile access to AnyConnect 2.3 and later versions. This license requires activation of one of the following licenses to specify the total number of SSL VPN sessions permitted: AnyConnect Essentials or AnyConnect Premium. Mobile Posture Support Enforcing remote access controls and gathering posture data from mobile devices requires an AnyConnect Mobile license and either an AnyConnect Essentials or AnyConnect Premium license to be installed on the ASA. Here is the functionality you receive based on the license you install. - AnyConnect Premium License Functionality - Enforce DAP policies on supported mobile devices based on DAP attributes and any other existing endpoint attributes. This includes allowing or denying remote access from a mobile device. - AnyConnect Essentials License Functionality - Enable or disable mobile device access on a per group basis and to configure that feature using ASDM. - Display information about connected mobile devices via CLI or ASDM without having the ability to enforce DAP policies or deny or allow remote access to those mobile devices.
AnyConnect Premium	Note This license is a legacy license. We recommend that you upgrade to the AnyConnect Plus or Apex license, which includes this license. AnyConnect Premium sessions include the following VPN types: - SSL VPN - Clientless SSL VPN - IPsec remote access VPN using IKEv2
AnyConnect Premium Shared	Note This license is a legacy license. We recommend that you upgrade to the AnyConnect Plus or Apex license, which replaces the shared license function. A shared license lets the ASA act as a shared license server for multiple client ASAs. The shared license pool is large, but the maximum number of sessions used by each individual ASA cannot exceed the maximum number listed for permanent licenses.
Botnet Traffic Filter	Requires a Strong Encryption (3DES/AES) License to download the dynamic database.
Encryption	The DES license cannot be disabled. If you have the 3DES license installed, DES is still available. To prevent the use of DES when you want to only use strong encryption, be sure to configure any relevant commands to use only strong encryption.

Table 3-13 License Notes (continued)

License	Notes
Interfaces of all types, Max.	The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every interface defined in the configuration counts against this limit.
IPS module	The IPS module license lets you run the IPS software module on the ASA. You also need the IPS signature subscription on the IPS side. See the following guidelines: • To buy the IPS signature subscription you need to have the ASA with IPS pre-installed (the part number must include “IPS”, for example ASA5515-IPS-K9); you cannot buy the IPS signature subscription for a non-IPS part number ASA. • For failover, you need the IPS signature subscription on both units; this subscription is not shared in failover, because it is not an ASA license. • For failover, the IPS signature subscription requires a unique IPS module license per unit. Like other ASA licenses, the IPS module license is technically shared in the failover cluster license. However, because of the IPS signature subscription requirements, you must buy a separate IPS module license for each unit in failover.
Other VPN	Other VPN sessions include the following VPN types: • IPsec remote access VPN using IKEv1 • IPsec site-to-site VPN using IKEv1 • IPsec site-to-site VPN using IKEv2 This license is included in the Base license.
Total VPN (sessions), combined all types	• Although the maximum VPN sessions add up to more than the maximum VPN AnyConnect and Other VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the ASA, so be sure to size your network appropriately. • If you start a clientless SSL VPN session and then start an AnyConnect client session from the portal, 1 session is used in total. However, if you start the AnyConnect client first (from a standalone client, for example) and then log into the clientless SSL VPN portal, then 2 sessions are used.

Table 3-13 License Notes (continued)

License	Notes
Total UC Proxy Sessions	Each TLS proxy session for Encrypted Voice Inspection is counted against the UC license limit. 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). Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used. You independently set the TLS proxy limit using the tls-proxy maximum-sessions command or in ASDM, using the Configuration > Firewall > Unified Communications > TLS Proxy pane. To view the limits of your model, enter the tls-proxy maximum-sessions ? command. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license. Note For license part numbers ending in “K8” (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in “K9” (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted. If you clear the configuration (using the clear configure all command, for example), then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the tls-proxy maximum-sessions command to raise the limit again (in ASDM, use the TLS Proxy pane). If you use failover and enter the write standby command or in ASDM, use File > Save Running Configuration to Standby Unit on the primary unit to force a configuration synchronization, the clear configure all command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning. You might also use SRTP encryption sessions for your connections: • For K8 licenses, SRTP sessions are limited to 250. • For K9 licenses, there is not limit. Note Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.
Virtual CPU	You must install a model license on the ASAv that sets the appropriate number of vCPUs. Until you install a license, throughput is limited to 100 Kbps so that you can perform preliminary connectivity tests. A model license is required for regular operation.
VLANs, Maximum	For an interface to count against the VLAN limit, you must assign a VLAN to it.
VPN Load Balancing	VPN load balancing requires a Strong Encryption (3DES/AES) License.

About PAK 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.

- [Preinstalled License, page 3-19](#)
- [Permanent License, page 3-19](#)
- [Time-Based Licenses, page 3-19](#)
- [Shared AnyConnect Premium Licenses \(AnyConnect 3 and Earlier\), page 3-22](#)
- [Failover or ASA Cluster Licenses, page 3-22](#)
- [No Payload Encryption Models, page 3-26](#)
- [Licenses FAQ, page 3-26](#)

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.

Related Topics

[Monitoring PAK Licenses, page 3-35](#)

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.

Related Topics

[How Permanent and Time-Based Licenses Combine, page 3-20](#)

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 order a Botnet Traffic Filter time-based license that is valid for 1 year.

- [Time-Based License Activation Guidelines, page 3-20](#)
- [How the Time-Based License Timer Works, page 3-20](#)
- [How Permanent and Time-Based Licenses Combine, page 3-20](#)
- [Stacking Time-Based Licenses, page 3-21](#)
- [Time-Based License Expiration, page 3-22](#)

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 3000-session Unified Communications license, and a 2000-session Unified Communications 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 Unified Communications license, you cannot also activate a standalone time-based 2000-session Unified Communications 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 stops counting down. The time-based license only counts down when the ASA is running. The system clock setting does not affect the license; only ASA uptime counts towards the license duration.

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. The following table 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 3-14 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.

Table 3-14 Time-Based License Combination Rules

Time-Based Feature	Combined License Rule
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.

Related Topics

[Monitoring PAK Licenses, page 3-35](#)

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 that 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 Unified Communications 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 Unified Communications license vs. a 2000-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.

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.

Related Topics

- [Activate or Deactivate Keys, page 3-29](#)
- [Time-Based License Expiration, page 3-22](#)

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.

For example, you have a time-based 2000-session Unified Communications license (active), a time-based 1000-session Unified Communications license (inactive), and a permanent 500-session Unified Communications license. While the 2000-session license expires, the ASA activates the 1000-session license. After the 1000-session license expires, the ASA uses the 500-session permanent license.

Related Topics

[Activate or Deactivate Keys, page 3-29](#)

Shared AnyConnect Premium Licenses (AnyConnect 3 and Earlier)



Note

The shared license feature on the ASA is not supported with AnyConnect 4 and later licensing. AnyConnect licenses are shared and no longer require a shared server or participant license.

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.

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.

- [Failover License Requirements and Exceptions, page 3-22](#)
- [ASA Cluster License Requirements and Exceptions, page 3-23](#)
- [How Failover or ASA Cluster Licenses Combine, page 3-24](#)
- [Loss of Communication Between Failover or ASA Cluster Units, page 3-25](#)
- [Upgrading Failover Pairs, page 3-25](#)

Failover License Requirements and Exceptions

Failover units do not require the same license on each unit. 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. There are some exceptions to this rule. See the following table for precise licensing requirements for failover.

Model	License Requirement
ASA 5506-X series	- Active/Standby—Security Plus License. - Active/Active—No Support. Note Each unit must have the same encryption license.
ASA 5512-X through ASA 5555-X	- ASA 5512-X—Security Plus License. - Other models—Base License. Note Each unit must have the same encryption license; each unit must have the same 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.
ASAv	- Active/Standby—Standard and Premium Licenses. - Active/Active—No Support. Note The standby unit requires the same model license as the primary unit; Each unit must have the same encryption license.
All other models	Base License. Note Each unit must have the same encryption license.

**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.

There are exceptions to this rule. See the following table for precise licensing requirements for clustering.

Model	License Requirement
ASA 5585-X	Cluster License. Note Each unit must have the same encryption license; each unit must have the same 10 GE I/O/Security Plus license (ASA 5585-X with SSP-10 and -20).
ASA 5512-X	Security Plus license. Note Each unit must have the same encryption license.
ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X	Base License. Note Each unit must have the same encryption license.
All other models	No support.

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 5525-Xs 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 that 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 5545-X 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.

Related Topics

[Monitoring PAK Licenses, page 3-35](#)

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, then you can fail over to the other unit while you reload. If both units require reloading, then you can reload them separately so that you have no downtime.

Related Topics

[Table 3-15 on page 3-29](#)

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, VPN and Unified Communications licenses will not be listed.

Related Topics

[Monitoring PAK Licenses, page 3-35](#)

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 for PAK Licenses

Context Mode Guidelines

In multiple context mode, apply the activation key in the system execution space.

Failover Guidelines

See [Failover or ASA Cluster Licenses, page 3-22](#).

Model Guidelines

- Smart Licensing is supported on the ASAv only.
- Shared licenses are not supported on the ASAv, ASA 5506-X, ASA 5508-X, and ASA 5516-X.

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

- 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, if you install the AnyConnect Essentials license (if it is available for your model), it is used instead of the above licenses. You can disable the AnyConnect Essentials license in the configuration to restore use of the other licenses using the **Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials** pane.

Configure PAK Licenses

This section describes how to obtain an activation key and how to active it. You can also deactivate a key.

- [Obtain an Activation Key, page 3-28](#)
- [Activate or Deactivate Keys, page 3-29](#)

Obtain 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.

Procedure

-
Step 1	Obtain the serial number for your ASA by choosing Configuration > Device Management > Licensing > Activation Key (in multiple context mode, view the serial number in the System execution space).
Step 2	If you are not already registered with Cisco.com, create an account.
Step 3	Go to the following licensing website: http://www.cisco.com/go/license
Step 4	Enter the following information, when prompted:

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

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

- Step 5** If you have additional Product Authorization Keys, repeat [Step 4](#) for each Product Authorization Key. After you enter all of the Product Authorization Keys, the final activation key provided includes all of the permanent features you registered.

Activate or Deactivate Keys

This section describes how to enter a new activation key, and how to activate and deactivate time-based keys.

Before You Begin

- 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. The following table lists the licenses that require reloading.

Table 3-15 Permanent License Reloading Requirements

Model	License Action Requiring Reload
All models	Downgrading the Encryption license.
ASAv	Downgrading the vCPU license.

Procedure

- Step 1** Choose **Configuration > Device Management**, and then choose the **Licensing > Activation Key** or **Licensing Activation Key** pane, depending on your model.
- Step 2** To enter a new activation key, either permanent or time-based, enter the new activation key in the **New Activation Key** field.
- The key is a five-element hexadecimal string with one space between each element. The leading 0x specifier is optional; all values are assumed to be hexadecimal. For example:
- ```
ASA0xd11b3d48 0xa80a4c0a 0x48e0fd1c 0xb0443480 0x843fc490
```
- You can install one permanent key, and multiple time-based keys. If you enter a new permanent key, it overwrites the already installed one. If you enter a new time-based key, then it is active by default and displays in the Time-based License Keys Installed table. The last time-based key that you activate for a given feature is the active one.
- Step 3** To activate or deactivate an installed time-based key, choose the key in the **Time-based License Keys Installed** table, and click either **Activate** or **Deactivate**.
- You can only have one time-based key active for each feature.

**Step 4** Click **Update Activation Key**.

Some permanent licenses require you to reload the ASA after entering the new activation key. You will be prompted to reload if it is required.

---

**Related Topics**

- [Time-Based Licenses, page 3-19](#)
- [Table 3-15 on page 3-29](#)

## Configure a Shared License (AnyConnect 3 and Earlier)

**Note**

The shared license feature on the ASA is not supported with AnyConnect 4 and later licensing. AnyConnect licenses are shared and no longer require a shared server or participant license.

---

This section describes how to configure the shared licensing server and participants.

- [About Shared Licenses, page 3-30](#)
- [Configure the Shared Licensing Server, page 3-34](#)
- [Configure the Shared Licensing Participant and the Optional Backup Server, page 3-34](#)

## About Shared 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.

- [About the Shared Licensing Server and Participants, page 3-30](#)
- [Communication Issues Between Participant and Server, page 3-31](#)
- [About the Shared Licensing Backup Server, page 3-32](#)
- [Failover and Shared Licenses, page 3-32](#)
- [Maximum Number of Participants, page 3-34](#)

## 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.

## 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.

- [Failover and Shared License Servers, page 3-32](#)
- [Failover and Shared License Participants, page 3-33](#)

### Failover and Shared License Servers

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

**Note**


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

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

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

For example, you have a network with 2 failover pairs. Pair #1 includes the main licensing server. Pair #2 includes the backup server. When the primary unit from Pair #1 goes down, the standby unit immediately becomes the new main licensing server. The backup server from Pair #2 never gets used. Only if both units in Pair #1 go down does the backup server in Pair #2 come into use as the shared licensing server. If Pair #1 remains down, and the primary unit in Pair #2 goes down, then the standby unit in Pair #2 comes into use as the shared licensing server (see [Figure 3-1](#)).

**Figure 3-1 Failover and Shared License Servers**



The standby backup server shares the same operating limits as the primary backup server; if the standby unit becomes active, it continues counting down where the primary unit left off.

#### Related Topics

[About the Shared Licensing Backup Server, page 3-32](#)

#### 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.

## Configure the Shared Licensing Server

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

### Before You Begin

The server must have a shared licensing server key.

### Procedure

- 
- Step 1** Choose the **Configuration > Device Management > Licenses > Shared SSL VPN Licenses** pane.
  - Step 2** In the **Shared Secret** field, enter the shared secret as a string between 4 and 128 ASCII characters.  
Any participant with this secret can use the license server.
  - Step 3** (Optional) In the **TCP IP Port** field, enter the port on which the server listens for SSL connections from participants, between 1 and 65535.  
The default is TCP port 50554.
  - Step 4** (Optional) In the **Refresh interval** field, enter the refresh interval between 10 and 300 seconds.  
This value is provided to participants to set how often they should communicate with the server. The default is 30 seconds.
  - Step 5** In the **Interfaces that serve shared licenses** area, check the **Shares Licenses** check box for any interfaces on which participants contact the server.
  - Step 6** (Optional) To identify a backup server, in the **Optional backup shared SSL VPN license server** area:
    - a.** In the **Backup server IP address** field, enter the backup server IP address.
    - b.** In the **Primary backup server serial number** field, enter the backup server serial number.
    - c.** If the backup server is part of a failover pair, identify the standby unit serial number in the **Secondary backup server serial number** field.

You can only identify 1 backup server and its optional standby unit.
  - Step 7** Click **Apply**.
- 

## Configure the Shared Licensing Participant and the Optional Backup Server

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

### Before You Begin

The participant must have a shared licensing participant key.



### Procedure

- 
- Step 1** Choose the **Configuration > Device Management > Licenses > Shared SSL VPN Licenses** pane.
- Step 2** In the Shared Secret field, enter the shared secret as a string between 4 and 128 ASCII characters.
- Step 3** (Optional) In the TCP IP Port field, enter the port on which to communicate with the server using SSL, between 1 and 65535.
- The default is TCP port 50554.
- Step 4** (Optional) To identify the participant as the backup server, in the Select backup role of participant area:
- Click the **Backup Server** radio button.
  - Check the **Shares Licenses** check box for any interfaces on which participants contact the backup server.
- Step 5** Click **Apply**.
- 

## Monitoring PAK Licenses

This section describes how to view license information.

- [Viewing Your Current License, page 3-35](#)
- [Monitoring the Shared License, page 3-36](#)

## 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.

### Before You Begin

If you have a No Payload Encryption model, then you view the license, VPN and Unified Communications licenses will not be listed. See [No Payload Encryption Models, page 3-26](#) for more information.

### Procedure

- 
- Step 1** To view the running license, which is a combination of the permanent license and any active time-based licenses, choose the **Configuration > Device Management > Licensing > Activation Key** pane and view the Running Licenses area.
- In multiple context mode, view the activation key in the System execution space by choosing the **Configuration > Device Management > Activation Key** pane.
- For a failover pair, the running license shown is the combined license from the primary and secondary units. See [How Failover or ASA Cluster Licenses Combine, page 3-24](#) for more information. For time-based licenses with numerical values (the duration is not combined), the License Duration column displays the shortest time-based license from either the primary or secondary unit; when that license expires, the license duration from the other unit displays.

- Step 2** (Optional) To view time-based license details, such as the features included in the license and the duration, in the Time-Based License Keys Installed area, choose a license key, and then click **Show License Details**.
- Step 3** (Optional) For a failover unit, to view the license installed on this unit (and not the combined license from both primary and secondary units), in the Running Licenses area, click **Show information of license specifically purchased for this device alone**.

## Monitoring the Shared License

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

## History for PAK Licenses

| Feature Name                                              | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Increased Connections and VLANs                           | 7.0(5)            | Increased the following limits: <ul style="list-style-type: none"> <li>ASA5510 Base license connections from 32000 to 5000; VLANs from 0 to 10.</li> <li>ASA5510 Security Plus license connections from 64000 to 130000; VLANs from 10 to 25.</li> <li>ASA5520 connections from 130000 to 280000; VLANs from 25 to 100.</li> <li>ASA5540 connections from 280000 to 400000; VLANs from 100 to 200.</li> </ul> |
| 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.                                                                                                                                                                                                                                                              |

| Feature Name                                                    | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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><b>Note</b> The interface names remain Ethernet 0/0 and Ethernet 0/1.</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 antispymware 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)            | VPN load balancing is now supported on the ASA 5510 Security Plus license.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| AnyConnect for Mobile License                                   | 8.0(3)            | The AnyConnect for Mobile license was introduced. It lets Windows mobile devices connect to the ASA using the AnyConnect client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Time-based Licenses                                             | 8.0(4)/8.1(2)     | Support for time-based licenses was introduced.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Feature Name                                                                       | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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><b>Note</b> 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 Configuration &gt; Remote Access VPN &gt; Network (Client) Access &gt; Advanced &gt; AnyConnect Essentials pane.</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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Feature Name                                                                             | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10 GE I/O license for the ASA 5585-X with SSP-20                                         | 8.2(3)            | <p>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.</p> <p><b>Note</b> The ASA 5585-X is not supported in 8.3(x).</p>                                                                                                                                                                                                                                                                                    |
| 10 GE I/O license for the ASA 5585-X with SSP-10                                         | 8.2(4)            | <p>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.</p> <p><b>Note</b> The ASA 5585-X is not supported in 8.3(x).</p>                                                                                                                                                                                                                                                                                    |
| Non-identical failover licenses                                                          | 8.3(1)            | <p>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.</p> <p>We modified the following screen: Configuration &gt; Device Management &gt; Licensing &gt; Activation Key.</p>                                                                                                                                                                                                                                                           |
| Stackable time-based licenses                                                            | 8.3(1)            | <p>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 that you do not have to worry about the license expiring or about losing time on your licenses because you installed the new one early.</p> |
| Intercompany Media Engine License                                                        | 8.3(1)            | The IME license was introduced.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Time-based licenses based on Uptime                                                      | 8.3(1)            | Time-based licenses now count down according to the total uptime of the ASA; the system clock does not affect the license.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Multiple time-based licenses active at the same time                                     | 8.3(1)            | <p>You can now install multiple time-based licenses, and have one license per feature active at a time.</p> <p>The following screen was modified: Configuration &gt; Device Management &gt; Licensing &gt; Activation Key.</p>                                                                                                                                                                                                                                                                                                                  |
| Discrete activation and deactivation of time-based licenses.                             | 8.3(1)            | <p>You can now activate or deactivate time-based licenses using a command.</p> <p>We modified the following screen: Configuration &gt; Device Management &gt; Licensing &gt; Activation Key.</p>                                                                                                                                                                                                                                                                                                                                                |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Feature Name                                                             | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No Payload Encryption image for export                                   | 8.3(2)            | <p>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.</p> <p><b>Note</b> 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.</p>                                        |
| 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)            | <p>We increased the firewall connection limits:</p> <ul style="list-style-type: none"> <li>ASA 5580-20—1,000,000 to 2,000,000.</li> <li>ASA 5580-40—2,000,000 to 4,000,000.</li> <li>ASA 5585-X with SSP-10: 750,000 to 1,000,000.</li> <li>ASA 5585-X with SSP-20: 1,000,000 to 2,000,000.</li> <li>ASA 5585-X with SSP-40: 2,000,000 to 4,000,000.</li> <li>ASA 5585-X with SSP-60: 2,000,000 to 10,000,000.</li> </ul>       |
| 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)            | <p>IPsec remote access VPN using IKEv2 was added to the AnyConnect Essentials and AnyConnect Premium licenses.</p> <p><b>Note</b> The following limitation exists in our support for IKEv2 on the ASA:<br/>We currently do not support duplicate security associations.</p> <p>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.</p> |
| 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.                                                                                                                                                                                                                     |

| Feature Name                                                                                                            | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.                                                                                                                                                                                                             |
| Support for 16 cluster members for the ASA 5585-X                                                                       | 9.2(1)            | The ASA 5585-X now supports 16-unit clusters.                                                                                                                                                                                                                                                                                                                                                                                         |
| ASAv4 and ASAv30 Standard and Premium model licenses introduced                                                         | 9.2(1)            | The ASAv was introduced with a simple licensing scheme: ASAv4 and ASAv30 permanent licenses in Standard or Premium levels. No add-on licenses are available.                                                                                                                                                                                                                                                                          |







## Smart Software Licensing for the ASAv

---

Cisco Smart Software Licensing lets you purchase and manage a pool of licenses centrally. Unlike product authorization key (PAK) licenses, smart licenses are not tied to a specific serial number. You can easily deploy or retire ASAvs without having to manage each unit's license key. Smart Software Licensing also lets you see your license usage and needs at a glance.

- [Supported Licenses, page 5-1](#)
- [About Smart Software Licensing, page 5-4](#)
- [Prerequisites for Smart Software Licensing, page 5-6](#)
- [Guidelines for Smart Software Licensing, page 5-6](#)
- [Defaults for Smart Software Licensing, page 5-6](#)
- [Configure Smart Software Licensing, page 5-7](#)
- [Manage Smart Software Licensing, page 5-8](#)
- [Monitoring Smart Software Licensing, page 5-9](#)
- [History for Smart Software Licensing, page 5-10](#)

## Supported Licenses

This section lists the license entitlements available for the ASAv.

- [ASAv5 and ASAv10, page 5-2](#)
- [ASAv30, page 5-3](#)
- [License Notes, page 5-4](#)

## ASAv5 and ASAv10

Table 5-1 ASAv5 and ASAv10 License Features

| Licenses                                                                                                                                                        | Standard License                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| <b>Firewall Licenses</b>                                                                                                                                        |                                     |
| Botnet Traffic Filter                                                                                                                                           | Supported                           |
| Firewall Conns, Concurrent                                                                                                                                      | 100,000                             |
| GTP/GPRS                                                                                                                                                        | Supported                           |
| Intercompany Media Eng.                                                                                                                                         | Supported                           |
| UC Phone Proxy Sessions,<br>Total UC Proxy Sessions                                                                                                             | 500                                 |
| <b>VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.</b> |                                     |
| AnyConnect Premium Peers<br>(maximum)                                                                                                                           | 250                                 |
|                                                                                                                                                                 | <i>Shared licenses: No Support.</i> |
| Adv. Endpoint Assessment                                                                                                                                        | Enabled                             |
| AnyConnect for Cisco VPN<br>Phone                                                                                                                               | Enabled                             |
| AnyConnect Essentials                                                                                                                                           | Disabled                            |
| AnyConnect for Mobile                                                                                                                                           | Enabled                             |
| <b>Other VPN Licenses</b>                                                                                                                                       |                                     |
| Total VPN Peers, combined<br>all types                                                                                                                          | 250                                 |
| Other VPN Peers                                                                                                                                                 | 250                                 |
| VPN Load Balancing                                                                                                                                              | Enabled                             |
| <b>General Licenses</b>                                                                                                                                         |                                     |
| Throughput Level                                                                                                                                                | ASAv5: 100 Mbps<br>ASAv10: 1 Gbps   |
| Encryption                                                                                                                                                      | Strong (3DES/AES)                   |
| Failover                                                                                                                                                        | Active/Standby                      |
| Interfaces of all types, Max.                                                                                                                                   | 716                                 |
| Security Contexts                                                                                                                                               | No support                          |
| Clustering                                                                                                                                                      | No support                          |
| VLANs, Maximum                                                                                                                                                  | 50                                  |
| RAM, vCPUs, vCPU<br>Frequency Limit                                                                                                                             | 2 GB, 1 vCPU, 5000 MHz              |

## ASAv30

Table 5-2 ASAv30 License Features

| Licenses                                                                                                                                                        | Standard License                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Firewall Licenses</b>                                                                                                                                        |                                                                                                                                                                                                                                                                                            |
| Botnet Traffic Filter                                                                                                                                           | Supported                                                                                                                                                                                                                                                                                  |
| Firewall Conns, Concurrent                                                                                                                                      | 500,000                                                                                                                                                                                                                                                                                    |
| GTP/GPRS                                                                                                                                                        | Supported                                                                                                                                                                                                                                                                                  |
| Intercompany Media Eng.                                                                                                                                         | Supported                                                                                                                                                                                                                                                                                  |
| UC Phone Proxy Sessions,<br>Total UC Proxy Sessions                                                                                                             | 1000                                                                                                                                                                                                                                                                                       |
| <b>VPN Licenses require an AnyConnect Plus or Apex license, available separately. See the following maximum values when you purchase an AnyConnect license.</b> |                                                                                                                                                                                                                                                                                            |
| AnyConnect Premium Peers<br>(maximum)                                                                                                                           | 750                                                                                                                                                                                                                                                                                        |
|                                                                                                                                                                 | <i>Shared licenses: No Support.</i>                                                                                                                                                                                                                                                        |
| Adv. Endpoint Assessment                                                                                                                                        | Enabled                                                                                                                                                                                                                                                                                    |
| AnyConnect for Cisco VPN<br>Phone                                                                                                                               | Enabled                                                                                                                                                                                                                                                                                    |
| AnyConnect Essentials                                                                                                                                           | Disabled                                                                                                                                                                                                                                                                                   |
| AnyConnect for Mobile                                                                                                                                           | Enabled                                                                                                                                                                                                                                                                                    |
| <b>Other VPN Licenses</b>                                                                                                                                       |                                                                                                                                                                                                                                                                                            |
| Total VPN Peers, combined<br>all types                                                                                                                          | 750                                                                                                                                                                                                                                                                                        |
| Other VPN Peers                                                                                                                                                 | 750                                                                                                                                                                                                                                                                                        |
| VPN Load Balancing                                                                                                                                              | Enabled                                                                                                                                                                                                                                                                                    |
| <b>General Licenses</b>                                                                                                                                         |                                                                                                                                                                                                                                                                                            |
| Throughput Level                                                                                                                                                | 2 Gbps                                                                                                                                                                                                                                                                                     |
| Encryption                                                                                                                                                      | Strong (3DES/AES)                                                                                                                                                                                                                                                                          |
| Failover                                                                                                                                                        | Active/Standby                                                                                                                                                                                                                                                                             |
| Interfaces of all types, Max.                                                                                                                                   | 1316                                                                                                                                                                                                                                                                                       |
| Security Contexts                                                                                                                                               | No support                                                                                                                                                                                                                                                                                 |
| Clustering                                                                                                                                                      | No support                                                                                                                                                                                                                                                                                 |
| VLANs, Maximum                                                                                                                                                  | 200                                                                                                                                                                                                                                                                                        |
| RAM, vCPUs, vCPU<br>Frequency Limit                                                                                                                             | 8 GB, 4 vCPUs, 20000 MHz                                                                                                                                                                                                                                                                   |
|                                                                                                                                                                 | <b>Note</b> If you choose to deploy 2 or 3 vCPUs, then see the following values:<br><br>2 vCPUs—4 GB RAM, vCPU Frequency Limit of 10000 MHz, 250,000 concurrent firewall connections.<br><br>3 vCPUs—4 GB RAM, vCPU Frequency Limit of 15000 MHz, 350,000 concurrent firewall connections. |

**License Notes**

The following table includes additional information about licenses.

**Table 5-3 License Notes**

| License                                  | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AnyConnect Premium                       | VPN Licenses require an AnyConnect Plus or Apex license, available separately.<br>AnyConnect Premium sessions include the following VPN types: <ul style="list-style-type: none"> <li>• SSL VPN</li> <li>• Clientless SSL VPN</li> <li>• IPsec remote access VPN using IKEv2</li> </ul>                                                                                                                                                                                                                                                                                                                                                                         |
| Encryption                               | The DES license cannot be disabled. Although the 3DES license is 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.                                                                                                                                                                                                                                                                                                                                                                                                            |
| Interfaces of all types, Max.            | The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every <b>interface</b> defined in the configuration counts against this limit.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Other VPN                                | Other VPN sessions include the following VPN types: <ul style="list-style-type: none"> <li>• IPsec remote access VPN using IKEv1</li> <li>• IPsec site-to-site VPN using IKEv1</li> <li>• IPsec site-to-site VPN using IKEv2</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Total VPN (sessions), combined all types | <ul style="list-style-type: none"> <li>• 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.</li> <li>• 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.</li> </ul> |
| VLANs, Maximum                           | For an interface to count against the VLAN limit, you must assign a VLAN to it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## About Smart Software Licensing

This section describes how Smart Software Licensing work with the ASAv.

- [Smart Software Manager and Accounts, page 5-5](#)
- [Licenses and Devices Managed per Virtual Account, page 5-5](#)
- [Device Registration and Tokens, page 5-5](#)
- [Periodic Communication with the License Authority, page 5-5](#)
- [Out-of-Compliance State, page 5-5](#)

## Smart Software Manager and Accounts

When you purchase 1 or more licenses for the ASAv, you manage them in the Cisco Smart Software Manager:

<http://tools.cisco.com/rhodu/index>

The Smart Software Manager lets you create a master account for your organization.

By default, your licenses are assigned to the *Default Virtual Account* under your master account. As the account administrator, you can optionally create additional virtual accounts; for example, you can create accounts for regions, departments, or subsidiaries. Multiple virtual accounts let you more easily manage large numbers of licenses and devices.

## Licenses and Devices Managed per Virtual Account

Licenses and devices are managed per virtual account: only that virtual account's ASAvs can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer ASAvs between virtual accounts.

## Device Registration and Tokens

For each virtual account, you can create a registration token. This token is valid for 30 days by default. Enter this token ID plus entitlement levels when you deploy each ASAv, or when you register an existing ASAv. You can create a new token if an existing token is expired.

At startup after deployment, or after you manually configure these parameters on an existing ASAv, the ASAv registers with the Cisco License Authority. When the ASAv registers with the token, the License Authority issues an ID certificate for communication between the device and the License Authority. This certificate is valid for 1 year, although it will be renewed every 6 months.

## Periodic Communication with the License Authority

The ASAv communicates with the License Authority every 30 days. If you make changes in the Smart Software Manager, you can refresh the authorization on the ASAv so the change takes place immediately. Or you can wait for the ASAv to communicate as scheduled.

You can optionally configure an HTTP proxy. Your ASAvs must have Internet access either directly or through an HTTP proxy at least every 90 days. Normal license communication occurs every 30 days, but with the grace period, your ASAv will operate for up to 90 days without calling home. You must contact the Licensing Authority before 90 days have passed.



**Note**

Offline licensing is not supported.

## Out-of-Compliance State

The ASAv can become out of compliance in the following situations:

- Over-utilization—When the ASAv uses unavailable licenses.

- License expiration—When a time-based license expires.
- Lack of communication—When the ASAv cannot reach the Licensing Authority for re-authorization.

After 90 days of reauthorization attempts, the ASAv will be severely rate-limited until you are able to successfully reauthorize.

## Smart Call Home Infrastructure

By default, a Smart Call Home profile exists in the configuration called “License”. This profile specifies the URL for the Licensing Authority. You cannot remove this profile. Note that the only configurable option for the License profile is the destination address URL for the License Authority. Unless directed by Cisco TAC, you should not change the License Authority URL.

You cannot disable Smart Call Home for Smart Software Licensing; for example, even if you disable Smart Call Home using the `no service call-home` command, Smart Software Licensing is not disabled.

Other Smart Call Home functions are not turned on unless you specifically configure them.

## Prerequisites for Smart Software Licensing

- Create a master account on the Cisco Smart Software Manager:  
<http://tools.cisco.com/rhodui/index>
- Purchase 1 or more ASAv licenses from Cisco Software Central.
- Ensure Internet access or HTTP proxy access from the ASAv, so the ASAv can contact the Licensing Authority. Offline licensing is not supported.
- Configure a DNS server so the ASA can resolve the name of the licensing authority server. See [Configure the DNS Server, page 18-10](#).

## Guidelines for Smart Software Licensing

### Failover

You must deploy both units with the same model license.

### Additional Guidelines

You cannot use PAK-based licensing with the ASAv. Only Smart Software Licensing is supported. If you upgrade an existing PAK-licensed ASAv, then the previously installed activation key will be ignored, but retained on the device. If you downgrade the ASAv, the activation key will be reinstated.

## Defaults for Smart Software Licensing

- The ASAv default configuration includes a Smart Call Home profile called “License” that specifies the URL for the Licensing Authority.
- When you deploy the ASAv, you set the feature tier and throughput level. Only the standard level is available at this time.

- Also during deployment, you can optionally configure an HTTP proxy.

## Configure Smart Software Licensing

When you deploy the ASAv, the device registers with the License Authority, and Smart Software Licensing is enabled according to the values you enter at deployment. If you need to change your license entitlement or configure Smart Software Licensing for an existing ASAv, perform the following tasks:

- 
- Step 1**    [\(Optional\) Configure the HTTP Proxy, page 5-7.](#)
  - Step 2**    [Set the Smart License Entitlement, page 5-7.](#)
  - Step 3**    [Register the ASAv with the License Authority, page 5-8.](#)
- 

### (Optional) Configure the HTTP Proxy

If your network uses an HTTP proxy for Internet access, you must configure the proxy address for Smart Software Licensing. This proxy is also used for Smart Call Home in general.

#### Procedure

- 
- Step 1**    Choose **Configuration > Device Management > Smart Call-Home**.
  - Step 2**    Check **Enable HTTP Proxy**.
  - Step 3**    Enter the proxy IP address and port in the **Proxy server** and **Proxy port** fields. For example, enter port 443 for an HTTPS server.
  - Step 4**    Click **Apply**.
- 

## Set the Smart License Entitlement

To request the license entitlement, perform the following procedure.

#### Procedure

- 
- Step 1**    Choose **Configuration > Device Management > Licensing > Smart Licensing**.
  - Step 2**    Check **Enable Smart license configuration**.
  - Step 3**    From the **Feature Tier** drop-down menu, choose **Standard**.  
Only the standard tier is available.
  - Step 4**    From the **Throughput Level** drop-down menu, choose **100M**, **1G**, or **2G**.

**Step 5** Click **Apply**.

---

## Register the ASAv with the License Authority

When you register the ASAv, the License Authority issues an ID certificate for communication between the ASAv and the License Authority. It also assigns the ASAv to the appropriate virtual account. Normally, this procedure is a one-time instance. However, you might need to later re-register the ASAv if the ID certificate expires because of a communication problem, for example.

### Procedure

---

- Step 1** In the Smart Software Manager, request and copy a registration token for the virtual account to which you want to add this ASAv.
- Step 2** Choose **Configuration > Device Management > Licensing > Smart Licensing**.
- Step 3** Click **Register**.
- Step 4** Enter the registration token in the **ID Token** field.
- Step 5** (Optional) Click the **Force registration** checkbox to register an ASAv that is already registered, but that might be out of sync with the License Authority.
- For example, use **Force registration** if the ASAv was accidentally removed from the Smart Software Manager.
- Step 6** Click **Register**.
- The ASAv attempts to register with the License Authority and request authorization for the configured license entitlements.
- 

## Manage Smart Software Licensing

You might need to deregister an ASAv from your account or manually renew the ID certificate or license entitlement.

- [Deregister the ASAv, page 5-8](#)
- [Renew the ID Certificate or License Entitlement, page 5-9](#)

## Deregister the ASAv

Deregistering the ASAv removes the ASAv from your account. All license entitlements and certificates on the ASAv are removed. You might want to deregister to free up a license for a new ASAv. Alternatively, you can remove the ASAv from the Smart Software Manager.

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Licensing > Smart Licensing**.



**Step 2** Click **Unregister**.

---

## Renew the ID Certificate or License Entitlement

By default, the ID certificate is automatically renewed every 6 months, and the license entitlement is renewed every 30 days. You might want to manually renew the registration for either of these items if you have a limited window for Internet access, or if you make any licensing changes in the Smart Software Manager, for example.

### Procedure

---

**Step 1** Choose **Configuration > Device Management > Licensing > Smart Licensing**.

**Step 2** To renew the ID certificate, click **Renew ID Certificate**.

**Step 3** To renew the license entitlement, click **Renew Authorization**.

---

## Monitoring Smart Software Licensing

You can monitor the license features, status, and certificate, as well as enable debug messages.

- [Viewing Your Current License, page 5-9](#)
- [Viewing Smart License Status, page 5-9](#)

### Viewing Your Current License

See the followingscreen for viewing your license:

- **Configuration > Device Management > Licensing > Smart Licensing** pane and view the **Effective Running Licenses** area.

### Viewing Smart License Status

See the followingscreens for viewing license status:

- **Monitoring > Properties > Smart License**  
Displays the state of Smart Software Licensing, Smart Agent version, UDI information, Smart Agent state, global compliance status, the entitlements status, licensing certificate information and schedule Smart Agent tasks.
- **Configuration > Device Management > Licensing > Smart Licensing > Registration Status**  
Displays the current Smart License registration status.

# History for Smart Software Licensing

| Feature Name                               | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cisco Smart Software Licensing for the ASA | 9.3(2)            | <p>Smart Software Licensing lets you purchase and manage a pool of licenses. Unlike PAK licenses, smart licenses are not tied to a specific serial number. You can easily deploy or retire ASAs without having to manage each unit's license key. Smart Software Licensing also lets you see your license usage and needs at a glance.</p> <p>We introduced or modified the following screens:</p> <p><b>Configuration &gt; Device Management &gt; Licensing &gt; Smart License</b></p> <p><b>Configuration &gt; Device Management &gt; Smart Call-Home</b></p> <p><b>Monitoring &gt; Properties &gt; Smart License</b></p> |



## Transparent or Routed Firewall Mode

---

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.

- [About the Firewall Mode, page 6-1](#)
- [Default Settings, page 6-7](#)
- [Guidelines for Firewall Mode, page 6-7](#)
- [Set the Firewall Mode \(Single Mode\), page 6-9](#)
- [Configure ARP Inspection for the Transparent Firewall, page 6-10](#)
- [Customize the MAC Address Table for the Transparent Firewall, page 6-11](#)
- [Examples for Firewall Mode, page 6-13](#)
- [History for the Firewall Mode, page 6-24](#)

### About the Firewall Mode

- [About Routed Firewall Mode, page 6-1](#)
- [About Transparent Firewall Mode, page 6-2](#)

### About Routed Firewall Mode

In routed mode, the Cisco 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.

## About Transparent Firewall Mode

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices.


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

## Using the Transparent Firewall in Your Network

The ASA connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

Figure 6-1 shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

**Figure 6-1** Transparent Firewall Network




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 6-2 shows two networks connected to the ASA, which has two bridge groups.

**Figure 6-2** Transparent Firewall Network with Two Bridge Groups



**Note**

Each bridge group requires a management IP address. The ASA uses this IP address as the source address for packets originating from the bridge group. The management IP address must be on the same subnet as the connected network. For another method of management, see [Management Interface, page 6-4](#).

The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.

## Management Interface

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 [Management Interface, page 11-2](#).

## Allowing Layer 3 Traffic

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



**Note**

Broadcast and multicast traffic can be passed using access rules. See the 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.

- IPv6 neighbor discovery and router solicitation packets are allowed through the transparent firewall in both directions without an ACL.
- For Layer 3 traffic traveling from a low to a high security interface, an extended ACL is required on the low security interface. See 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 [Transparent Firewall Mode Requirements for Failover, page 9-14](#) for more information.

## MAC Address vs. Route Lookups

When the ASA runs in transparent mode, the outgoing interface of a packet is determined by performing a MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for the following traffic types:

- Traffic originating on the ASA—For example, if your syslog server is located on a remote network, you must use a static route so the ASA can reach that subnet.
- Traffic that is at least one hop away from the ASA with NAT enabled—The ASA needs to perform a route lookup to find the next hop gateway; you need to add a static route on the ASA for the real host address.
- Voice over IP (VoIP) and DNS traffic with inspection enabled, and the endpoint is at least one hop away from the ASA—For example, if you use the transparent firewall between a CCM and an H.323 gateway, and there is a router between the transparent firewall and the H.323 gateway, then you need to add a static route on the ASA for the H.323 gateway for successful call completion. If you enable NAT for the inspected traffic, a static route is required to determine the egress interface for the real host address that is embedded in the packet. Affected applications include:
  - CTIQBE
  - DNS
  - GTP
  - H.323
  - MGCP
  - RTSP
  - SIP
  - Skinny (SCCP)

## ARP Inspection

By default, all ARP packets are allowed through the ASA. You can control the flow of ARP packets by enabling ARP inspection.

When you enable ARP inspection, the ASA compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the ASA drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the ASA to either forward the packet out all interfaces (flood), or to drop the packet.

**Note**

The dedicated management interface, if present, never floods packets even if this parameter is set to flood.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address.



The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

## MAC Address Table

The ASA learns and builds a MAC address table in a similar way as a normal bridge or switch: when a device sends a packet through the ASA, the ASA adds the MAC address to its table. The table associates the MAC address with the source interface so that the ASA knows to send any packets addressed to the device out the correct interface.

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.

## 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 for Firewall Mode

### 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 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 [Guidelines for Routed and Transparent Mode Interfaces, page 15-4](#) 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 [Set the Firewall Mode \(Single Mode\), page 6-9](#) for information about backing up your configuration file.
- If you download a text configuration to the ASA that changes the mode with the **firewall transparent** command, be sure to put the command at the top of the configuration; the ASA changes the mode as soon as it reads the command and then continues reading the configuration you downloaded. If the command appears later in the configuration, the ASA clears all the preceding lines in the configuration.

### Unsupported Features in Transparent Mode

[Table 6-1](#) lists the features are not supported in transparent mode.

**Table 6-1** *Unsupported Features in Transparent Mode*

| Feature                   | Description                                                                                                                                                                                                                                                                                                                                                |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dynamic DNS               | —                                                                                                                                                                                                                                                                                                                                                          |
| DHCP relay                | The transparent firewall can act as a DHCP server, but it does not support the DHCP relay commands. DHCP relay is not required because you can allow DHCP traffic to pass through using two extended ACLs: one that allows 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                       | —                                                                                                                                                                                                                                                                                                                                                          |

**Table 6-1**      *Unsupported Features in Transparent Mode*

| Feature                             | Description                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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              | —                                                                                                                                                                                                                                                                                                                              |

## Set the Firewall Mode (Single Mode)

This section describes how to change the firewall mode using the CLI. For single mode and for the currently connected context in multiple mode (typically the admin context), you cannot change the mode in ASDM. For other multiple mode contexts, you can set the mode in ASDM for each context; see [Configure a Security Context, page 8-18](#).

**Note**

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

### Prerequisites

When you change modes, the ASA clears the running configuration (see [Guidelines for Firewall Mode, page 6-7](#) for more information).

- If you already have a populated configuration, be sure to back up your configuration before changing the mode; you can use this backup for reference when creating your new configuration.
- Use the CLI at the console port to change the mode. If you use any other type of session, including the ASDM Command Line Interface tool or SSH, you will be disconnected when the configuration is cleared, and you will have to reconnect to the ASA using the console port in any case.
- Set the mode within the context.

### Procedure

**Note**

To set the firewall mode to transparent and also configure ASDM management access after the configuration is cleared, see [Configure ASDM Access, page 2-7](#).

**Step 1**      Set the firewall mode to transparent:

```
firewall transparent
```

Example:

```
ciscoasa(config)# firewall transparent
```

To change the mode to routed, enter the **no firewall transparent** command.

**Note**

You are not prompted to confirm the firewall mode change; the change occurs immediately.

## Configure ARP Inspection for the Transparent Firewall

To configure ARP Inspection, perform the following steps:

- Step 1** Add static ARP entries according to [Add a Static ARP Entry, page 6-10](#). ARP inspection compares ARP packets with static ARP entries in the ARP table, so static ARP entries are required for this feature.
- Step 2** Enable ARP inspection according to [Enable ARP Inspection, page 6-11](#).

## Add 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.

### Procedure

- Step 1** Choose the **Configuration > Device Management > Advanced > ARP > ARP Static Table** pane.
- Step 2** (Optional) Enter a value in the **ARP Timeout** field, to set the ARP timeout for *dynamic* ARP entries.  
  
This field sets the amount of time before the ASA rebuilds the ARP table, between 60 to 4294967 seconds. The default is 14400 seconds. Rebuilding the ARP table automatically updates new host information and removes old host information. You might want to reduce the timeout because the host information changes frequently.
- Step 3** (Optional; 8.4(5) only) To allow non-connected subnets, check the **Allow non-connected subnets** check box. The ASA ARP cache only contains entries from directly-connected subnets by default. You can enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries.

You may want to use this feature if you use:

- Secondary subnets.
- Proxy ARP on adjacent routes for traffic forwarding.

**Step 4** Click **Add**.

The **Add ARP Static Configuration** dialog box appears.

**Step 5** Choose the interface attached to the host network, from the **Interface** drop-down list.

**Step 6** Enter the IP address of the host, in the **IP Address** field.

**Step 7** Enter the MAC address of the host, in the **MAC Address** field; for example, 00e0.1e4e.3d8b.

**Step 8** Check the **Proxy ARP** check box, to perform proxy ARP for this address.

If the ASA receives an ARP request for the specified IP address, then it responds with the specified MAC address.

**Step 9** Click **OK**, and then **Apply**.

---

## Enable ARP Inspection

This section describes how to enable ARP inspection.

### Procedure

---

**Step 1** Choose the **Configuration > Device Management > Advanced > ARP > ARP Inspection** pane.

**Step 2** Choose the interface row on which you want to enable ARP inspection, and click **Edit**.

The **Edit ARP Inspection** dialog box appears.

**Step 3** Check the **Enable ARP Inspection** check box, to enable ARP inspection.

**Step 4** (Optional) Check the **Flood ARP Packets** check box, to flood non-matching ARP packets.

By default, packets that do not match any element of a static ARP entry are flooded out all interfaces except the originating interface. If there is a mismatch between the MAC address, the IP address, or the interface, then the ASA drops the packet.

If you uncheck this check box, all non-matching packets are dropped, which restricts ARP through the ASA to only static entries.



---

**Note** The Management 0/0 or 0/1 interface or subinterface, if present, never floods packets even if this parameter is set to flood.

---

**Step 5** Click **OK**, and then **Apply**.

---

## Customize the MAC Address Table for the Transparent Firewall

This section describes how you can customize the MAC address table.

- [Add a Static MAC Address, page 6-12](#)
- [Disable MAC Address Learning, page 6-12](#)

## Add 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 [Add a Static ARP Entry, page 6-10](#)), a static MAC address entry is automatically added to the MAC address table.

To add a static MAC address to the MAC address table, perform the following steps

### Procedure

- 
- Step 1** Choose the **Configuration > Device Setup > Bridging > MAC Address Table** pane.
  - Step 2** (Optional) Enter a value in the **Dynamic Entry Timeout** field, to set the time a MAC address entry stays in the MAC address table before timing out.  
This value is between 5 and 720 minutes (12 hours). 5 minutes is the default.
  - Step 3** Click **Add**.  
The **Add MAC Address Entry** dialog box appears.
  - Step 4** Choose the source interface associated with the MAC address, from the **Interface Name** drop-down list.
  - Step 5** Enter the MAC address, in the **MAC Address** field.
  - Step 6** Click **OK**, and then **Apply**.
- 

### Procedure

## Disable MAC Address Learning

By default, each interface automatically learns the MAC addresses of entering traffic, and the ASA adds corresponding entries to the MAC address table. You can disable MAC address learning if desired, however, unless you statically add MAC addresses to the table, no traffic can pass through the ASA.

To disable MAC address learning, perform the following steps:

### Procedure

- 
- Step 1** Choose the **Configuration > Device Setup > Bridging > MAC Learning** pane.
  - Step 2** To disable MAC learning, choose an interface row, and click **Disable**.

**Step 3** To reenable MAC learning, click **Enable**.

**Step 4** Click **Apply**.

---

## Examples for Firewall Mode

This section includes examples of how traffic moves through the ASA.

- [How Data Moves Through the ASA in Routed Firewall Mode, page 6-13](#)
- [How Data Moves Through the Transparent Firewall, page 6-19](#)

## How Data Moves Through the ASA in Routed Firewall Mode


This section describes how data moves through the ASA in routed firewall mode.

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

## An Inside User Visits a Web Server

Figure 6-3 shows an inside user accessing an outside web server.

**Figure 6-3**      *Inside to Outside*



The following steps describe how data moves through the ASA (see Figure 6-3):


1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).  
For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA translates the local source address (10.1.2.27) to the global address 209.165.201.10, which is on the outside interface subnet.  
The global address could be on any subnet, but routing is simplified when it is on the outside interface subnet.
4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. When www.example.com responds to the request, the packet goes through the ASA, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by untranslating the global destination address to the local user address, 10.1.2.27.
6. The ASA forwards the packet to the inside user.



## An Outside User Visits a Web Server on the DMZ

Figure 6-4 shows an outside user accessing the DMZ web server.

**Figure 6-4** *Outside to DMZ*




The following steps describe how data moves through the ASA (see Figure 6-4):

1. A user on the outside network requests a web page from the DMZ web server using the global destination address of 209.165.201.3, which is on the outside interface subnet.
2. The ASA receives the packet and untranslates the destination address to the local address 10.1.1.3.
3. Because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).  
For multiple context mode, the ASA first classifies the packet to a context.
4. The ASA then adds a session entry to the fast path and forwards the packet from the DMZ interface.
5. When the DMZ web server responds to the request, the packet goes through the ASA and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by translating the local source address to 209.165.201.3.
6. The ASA forwards the packet to the outside user.

## An Inside User Visits a Web Server on the DMZ

Figure 6-5 shows an inside user accessing the DMZ web server.

**Figure 6-5** Inside to DMZ




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

1. A user on the inside network requests a web page from the DMZ web server using the destination address of 10.1.1.3.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).  
For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA then records that a session is established and forwards the packet out of the DMZ interface.
4. When the DMZ web server responds to the request, the packet goes through the fast path, which lets the packet bypass the many lookups associated with a new connection.
5. The ASA forwards the packet to the inside user.

## An Outside User Attempts to Access an Inside Host

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

**Figure 6-6**      **Outside to Inside**

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


1. A user on the outside network attempts to reach an inside host (assuming the host has a routable IP address).  
If the inside network uses private addresses, no outside user can reach the inside network without NAT. The outside user might attempt to reach an inside user by using an existing NAT session.
2. The ASA receives the packet and because it is a new session, the ASA verifies if the packet is allowed according to the security policy (access lists, filters, AAA).
3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

## A DMZ User Attempts to Access an Inside Host

Figure 6-7 shows a user in the DMZ attempting to access the inside network.

**Figure 6-7** DMZ to Inside



The following steps describe how data moves through the ASA (see Figure 6-7):


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

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

## How Data Moves Through the Transparent Firewall

Figure 6-8 shows a typical transparent firewall implementation with an inside network that contains a public web server. The ASA has an access list so that the inside users can access Internet resources. Another access list lets the outside users access only the web server on the inside network.

**Figure 6-8** Typical Transparent Firewall Data Path




This section describes how data moves through the ASA.

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

## An Inside User Visits a Web Server

Figure 6-9 shows an inside user accessing an outside web server.

**Figure 6-9**      *Inside to Outside*



The following steps describe how data moves through the ASA (see Figure 6-9):

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

For multiple context mode, the ASA first classifies the packet to a context.

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


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

5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The ASA forwards the packet to the inside user.

## An Inside User Visits a Web Server Using NAT

Figure 6-10 shows an inside user accessing an outside web server.

**Figure 6-10** *Inside to Outside with NAT*



The following steps describe how data moves through the ASA (see Figure 6-10):

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

For multiple context mode, the ASA first classifies the packet according to a unique interface.

3. The ASA translates the real address (10.1.2.27) to the mapped address 209.165.201.10.  
Because the mapped address is not on the same network as the outside interface, then be sure the upstream router has a static route to the mapped network that points to the ASA.
4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 10.1.2.1.


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

6. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
7. The ASA performs NAT by untranslating the mapped address to the real address, 10.1.2.27.

## An Outside User Visits a Web Server on the Inside Network

Figure 6-11 shows an outside user accessing the inside web server.

**Figure 6-11**      *Outside to Inside*



92409

The following steps describe how data moves through the ASA (see Figure 6-11):

1. A user on the outside network requests a web page from the inside web server.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

For multiple context mode, the ASA first classifies the packet to a context.

3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.  
If the destination MAC address is not in the ASA table, the ASA attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.


5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The ASA forwards the packet to the outside user.



## An Outside User Attempts to Access an Inside Host

Figure 6-12 shows an outside user attempting to access a host on the inside network.

**Figure 6-12** *Outside to Inside*



The following steps describe how data moves through the ASA (see Figure 6-12):

1. A user on the outside network attempts to reach an inside host.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies if the packet is allowed according to the terms of the security policy (access lists, filters, AAA).  
For multiple context mode, the ASA first classifies the packet to a context.
3. The packet is denied because there is no access list permitting the outside host, and the ASA drops the packet.
4. If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

# History for the Firewall Mode

**Table 6-2** 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: <b>firewall transparent</b>, <b>show firewall</b>.</p> <p>You cannot set the firewall mode in ASDM; you must use the command-line interface.</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: <b>arp</b>, <b>arp-inspection</b>, and <b>show arp-inspection</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| MAC address table                  | 7.0(1)            | <p>Transparent firewall mode uses a MAC address table.</p> <p>We introduced the following commands: <b>mac-address-table static</b>, <b>mac-address-table aging-time</b>, <b>mac-learn disable</b>, and <b>show mac-address-table</b>.</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><b>Note</b> 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 modified or introduced the following screens:</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces &gt; Add/Edit Bridge Group Interface</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces &gt; Add/Edit Interface</p> |

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

| Feature Name                                           | Platform Releases | Feature Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ARP cache additions for non-connected subnets          | 8.4(5)/9.1(2)     | <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"> <li>• Secondary subnets.</li> <li>• Proxy ARP on adjacent routes for traffic forwarding.</li> </ul> <p>We modified the following screen: Configuration &gt; Device Management &gt; Advanced &gt; ARP &gt; ARP Static Table.</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: <b>firewall transparent</b>.</p> <p>For single mode, you cannot set the firewall mode in ASDM; you must use the command-line interface.</p> <p>For multiple mode, we modified the following screen: Configuration &gt; Context Management &gt; Security Contexts.</p>                                                                                                                                                                                                                                                                                   |
| Transparent mode bridge group maximum increased to 250 | 9.3(1)            | <p>The bridge group maximum was increased from 8 to 250 bridge groups. You can configure up to 250 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group.</p> <p>We modified the following screens:</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces &gt; Add/Edit Bridge Group Interface</p> <p>Configuration &gt; Device Setup &gt; Interface Settings &gt; Interfaces &gt; Add/Edit Interface</p>                                                                                                                                                                                                      |





# Startup Wizard

---

This chapter describes the ASDM Startup Wizard, which guides you through the initial configuration of the Cisco ASA and helps you define basic settings.

- [Accessing the Startup Wizard, page 7-1](#)
- [Guidelines for the Startup Wizard, page 7-1](#)
- [Startup Wizard Screens, page 7-1](#)
- [History for the Startup Wizard, page 7-5](#)

## Accessing the Startup Wizard

To access the Startup Wizard, choose one of the following options:

- **Wizards > Startup Wizard.**
- **Configuration > Device Setup > Startup Wizard**, then click **Launch Startup Wizard**.

## Guidelines for the Startup Wizard

### Context Mode Guidelines

The Startup Wizard is not supported in the system context.

## Startup Wizard Screens

The actual sequence of screens is determined by your specified configuration selections. Each screen is available for all modes or models unless otherwise noted.

## Starting Point or Welcome

- Click the **Modify existing configuration** radio button to change the existing configuration.
- Click the **Reset configuration to factory defaults** radio button to set the configuration to the factory default values.

- Check the **Configure the IP address of the management interface** check box to configure the IP address and subnet mask of the Management 0/0 interface to be different from the default value (192.168.1.1).

**Note**

If you reset the configuration to factory defaults, you cannot undo these changes by clicking **Cancel** or by closing this screen.

In multiple context mode, this screen does not include any parameters.

## Basic Configuration

Set the hostname, domain name, and enable password in this screen.

**Related Topics**

[Set the Hostname, Domain Name, and the Enable and Telnet Passwords, page 18-1](#)

## Interface Screens

The interface screens depend on the mode and model selected.

### Outside Interface Configuration (Routed Mode)

- Configure the IP address of the outside interface (the interface with the lowest security level).
- Configure the IPv6 address.

**Related Topics**

- [Configure General Routed Mode Interface Parameters, page 15-6](#)
- [Configure IPv6 Addressing, page 15-12](#)

### Outside Interface Configuration - PPPoE (Routed Mode, Single Mode)

Configure the PPPoE settings for the outside interface.

**Related Topics**

[Configure a Management Interface, page 15-11](#)

### Management IP Address Configuration (Transparent Mode)

For IPv4, a management IP address is required for each bridge group for both management traffic and for traffic to pass through the ASA. This screen sets the IP address for BVI 1.

**Related Topics**

[Configure Bridge Groups, page 15-9](#)

## Other Interfaces Configuration

Configure parameters for other interfaces.

### Related Topics

- [Configure General Routed Mode Interface Parameters, page 15-6](#)
- [Allow Same Security Level Communication, page 16-7](#)

## Static Routes

Configure static routes.

### Related Topics

[Configure a Static Route, page 22-4](#)

## DHCP Server

Configure the DHCP server.

### Related Topics

[Configure the DHCP Server, page 19-5](#)

## Address Translation (NAT/PAT)

Configures NAT or PAT for inside addresses (the interface with the highest security level) when accessing the outside (the interface with the lowest security level). See the firewall configuration guide for more information.

## Administrative Access

- Configure ASDM, Telnet, or SSH access.
- Check the **Enable HTTP server for HTTPS/ASDM access** check box to enable a secure connection to an HTTP server to access ASDM.
- Check the **Enable ASDM history metrics** check box.

### Related Topics

- [Configure ASA Access for ASDM, Telnet, or SSH, page 34-4](#)
- [Enable History Metrics, page 3-31](#)

## IPS Basic Configuration

In single context mode, use the Startup Wizard in ASDM to configure basic IPS network configuration. These settings are saved to the IPS configuration, not the ASA configuration. See the firewall configuration guide for more information.

## ASA CX Basic Configuration (ASA 5585-X)

You can use the Startup Wizard in ASDM to configure the ASA CX management address and Auth Proxy Port. These settings are saved to the ASA CX configuration, not the ASA configuration. You will also need to set additional network settings at the ASA CX CLI. See the firewall configuration guide for information about this screen.

## ASA FirePOWER Basic Configuration

You can use the Startup Wizard in ASDM to configure the ASA FirePOWER management address information and accept the end user license agreement (EULA). These settings are saved to the ASA FirePOWER configuration, not the ASA configuration. You will also need to configure some settings in the ASA FirePOWER CLI. For more information, see the chapter on the ASA FirePOWER module in the firewall configuration guide.

## Time Zone and Clock Configuration

Configure the clock parameters.

### Related Topics

[Set the Date and Time, page 18-7](#)

## Auto Update Server (Single Mode)

- Configure an auto update server by checking the **Enable Auto Update Server for ASA** check box.
- Check the **Enable Signature and Engine Updates from Cisco.com** check box if you have an IPS module. Set the following additional parameters:
  - Enter your Cisco.com username and password, then confirm the password.
  - Enter the start time in hh:mm:ss format, using a 24-hour clock.

### Related Topics

[Configure Auto Update, page 35-26](#)

## Startup Wizard Summary

This screen summarizes all of the configuration settings that you have made for the ASA.

- Click **Back** to change any of the settings in previous screens.
- Choose one of the following:
  - If you ran the Startup Wizard directly from a browser, when you click **Finish**, the configuration settings that you created through the wizard are sent to the ASA and saved in flash memory automatically.
  - If you ran the Startup Wizard from within ASDM, you must explicitly save the configuration in flash memory by choosing **File > Save Running Configuration to Flash**.



# History for the Startup Wizard

**Table 7-1** History for the Startup Wizard

| Feature Name                | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup Wizard              | 7.0(1)            | This wizard was introduced.<br>We introduced the <b>Wizards &gt; Startup Wizard</b> screen.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| ASA IPS Configuration       | 8.4(1)            | For the ASA IPS module, the <b>IPS Basic Configuration</b> screen was added to the startup wizard. Signature updates for the IPS module were also added to the <b>Auto Update</b> screen. The <b>Time Zone and Clock Configuration</b> screen was added to ensure the clock is set on the ASA; the IPS module gets its clock from the ASA.<br>We introduced or modified the following screens:<br><b>Wizards &gt; Startup Wizard &gt; IPS Basic Configuration</b><br><b>Wizards &gt; Startup Wizard &gt; Auto Update</b><br><b>Wizards &gt; Startup Wizard &gt; Time Zone and Clock Configuration</b> |
| ASA CX Configuration        | 9.1(1)            | For the ASA CX module, the <b>ASA CX Basic Configuration</b> screen was added to the startup wizard.<br>We introduced the following screens:<br><b>Wizards &gt; Startup Wizard &gt; ASA CX Basic Configuration</b>                                                                                                                                                                                                                                                                                                                                                                                    |
| ASA FirePOWER Configuration | 9.2(2.4)          | For the ASA FirePOWER module, the <b>ASA FirePOWER Basic Configuration</b> screen was added to the startup wizard.<br>We introduced the following screens:<br><b>Wizards &gt; Startup Wizard &gt; ASA FirePOWER Basic Configuration</b>                                                                                                                                                                                                                                                                                                                                                               |





## **PART 2**

### **High Availability and Scalability**





# Multiple Context Mode

---

This chapter describes how to configure multiple security contexts on the Cisco ASA.

- [About Security Contexts, page 8-1](#)
- [Licensing for Multiple Context Mode, page 8-13](#)
- [Guidelines for Multiple Context Mode, page 8-14](#)
- [Defaults for Multiple Context Mode, page 8-14](#)
- [Configure Multiple Contexts, page 8-15](#)
- [Change Between Contexts and the System Execution Space, page 8-22](#)
- [Manage Security Contexts, page 8-22](#)
- [Monitoring Security Contexts, page 8-26](#)
- [History for Multiple Context Mode, page 8-28](#)

## 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 [Guidelines for Multiple Context Mode, page 8-14](#).

This section provides an overview of security contexts.

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

## Common Uses for Security Contexts

You might want to use multiple security contexts in the following situations:

- You are a service provider and want to sell security services to many customers. By enabling multiple security contexts on the ASA, you can implement a cost-effective, space-saving solution that keeps all customer traffic separate and secure, and also eases configuration.
- You are a large enterprise or a college campus and want to keep departments completely separate.
- You are an enterprise that wants to provide distinct security policies to different departments.
- You have any network that requires more than one ASA.

## Context Configuration Files

This section describes how the ASA implements multiple context mode configurations.

- [Context Configurations, page 8-2](#)
- [System Configuration, page 8-2](#)
- [Admin Context Configuration, page 8-2](#)

## Context Configurations

For each context, the ASA includes a configuration that identifies the security policy, interfaces, and all the options you can configure on a standalone device. You can store context configurations in flash memory, or you can download them from a TFTP, FTP, or HTTP(S) server.

## System Configuration

The system administrator adds and manages contexts by configuring each context configuration location, allocated interfaces, and other context operating parameters in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the ASA. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the *admin context*. The system configuration does include a specialized failover interface for failover traffic only.

## Admin Context Configuration

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users. The admin context must reside on flash memory, and not remotely.

If your system is already in multiple context mode, or if you convert from single mode, the admin context is created automatically as a file on the internal flash memory called `admin.cfg`. This context is named “admin.” If you do not want to use `admin.cfg` as the admin context, you can change the admin context.

## How the ASA Classifies Packets

Each packet that enters the ASA must be classified, so that the ASA can determine to which context to send a packet.

- [Valid Classifier Criteria, page 8-3](#)
- [Classification Examples, page 8-4](#)

**Note**

If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each context.

### Valid Classifier Criteria

This section describes the criteria used by the classifier.

- [Unique Interfaces, page 8-3](#)
- [Unique MAC Addresses, page 8-3](#)
- [NAT Configuration, page 8-3](#)

**Note**

For management traffic destined for an interface, the interface IP address is used for classification.

The routing table is not used for packet classification.

### Unique Interfaces

If only one context is associated with the ingress interface, the ASA classifies the packet into that context. In transparent firewall mode, unique interfaces for contexts are required, so this method is used to classify packets at all times.

### Unique MAC Addresses

If multiple contexts share an interface, then the classifier uses unique MAC addresses assigned to the interface in each context. An upstream router cannot route directly to a context without unique MAC addresses. By default, auto-generation of MAC addresses is enabled. You can also set the MAC addresses manually when you configure each interface.


### NAT Configuration

If you disable use of unique MAC addresses, then the ASA uses the mapped addresses in your NAT configuration to classify packets. We recommend using MAC addresses instead of NAT, so that traffic classification can occur regardless of the completeness of the NAT configuration.

## Classification Examples

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

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




153367




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

**Figure 8-2** Incoming Traffic from Inside Networks



For transparent firewalls, you must use unique interfaces. [Figure 8-3](#) shows a packet destined to a host on the Context B inside network from the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

**Figure 8-3** Transparent Firewall Contexts



## Cascading Security Contexts

Placing a context directly in front of another context is called *cascading contexts*; the outside interface of one context is the same interface as the inside interface of another context. You might want to cascade contexts if you want to simplify the configuration of some contexts by configuring shared parameters in the top context.




### Note

Cascading contexts requires unique MAC addresses for each context interface (the default setting). Because of the limitations of classifying packets on shared interfaces without MAC addresses, we do not recommend using cascading contexts without unique MAC addresses.

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

**Figure 8-4 Cascading Contexts**



## Management Access to Security Contexts

The ASA provides system administrator access in multiple context mode as well as access for individual context administrators. The following sections describe logging in as a system administrator or as a context administrator:

- [System Administrator Access, page 8-7](#)
- [Context Administrator Access, page 8-8](#)

### System Administrator Access

You can access the ASA as a system administrator in two ways:

- Access the ASA console.  
From the console, you access the *system execution space*, which means that any commands you enter affect only the system configuration or the running of the system (for run-time commands).
- Access the admin context using Telnet, SSH, or ASDM.

See [Chapter 34, “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 34, “Management Access,”](#) to enable Telnet, SSH, and ASDM access and to configure management authentication.

## 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.

- [Resource Classes, page 8-8](#)
- [Resource Limits, page 8-8](#)
- [Default Class, page 8-9](#)
- [Use Oversubscribed Resources, page 8-10](#)
- [Use Unlimited Resources, page 8-11](#)

## Resource Classes

The ASA manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. To use the settings of a class, assign the context to the class when you define the context. All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to default. You can only assign a context to one resource class. The exception to this rule is that limits that are undefined in the member class are inherited from the default class; so in effect, a context could be a member of default plus another class.

## Resource Limits

You can set the limit for individual resources as a percentage (if there is a hard system limit) or as an absolute value.

For most resources, the ASA does not set aside a portion of the resources for each context assigned to the class; rather, the ASA sets the maximum limit for a context. If you oversubscribe resources, or allow some resources to be unlimited, a few contexts can “use up” those resources, potentially affecting service

to other contexts. The exception is VPN resource types, which you cannot oversubscribe, so the resources assigned to each context are guaranteed. To accommodate temporary bursts of VPN sessions beyond the amount assigned, the ASA supports a “burst” VPN resource type, which is equal to the remaining unassigned VPN sessions. The burst sessions *can* be oversubscribed, and are available to contexts on a first-come, first-served basis.

## Default Class

All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to the default class.


If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 2 percent limit for all concurrent connections, but no other limits, then all other limits are inherited from the default class. Conversely, if you create a class with a limit for all resources, the class uses no settings from the default class.

For most resources, the default class provides unlimited access to resources for all contexts, except for the following limits:

- Telnet sessions—5 sessions. (The maximum per context.)
- SSH sessions—5 sessions. (The maximum per context.)
- IPsec sessions—5 sessions. (The maximum per context.)
- MAC addresses—65,535 entries. (The maximum per context.)
- VPN site-to-site tunnels—0 sessions. (You must manually configure the class to allow any VPN sessions.)

Figure 8-5 shows the relationship between the default class and other classes. Contexts A and C belong to classes with some limits set; other limits are inherited from the default class. Context B inherits no limits from default because all limits are set in its class, the Gold class. Context D was not assigned to a class, and is by default a member of the default class.


**Figure 8-5 Resource Classes**



## Use Oversubscribed Resources

You can oversubscribe the ASA by assigning more than 100 percent of a resource across all contexts (with the exception of non-burst VPN resources). For example, you can set the Bronze class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended. (See Figure 8-6.)


**Figure 8-6 Resource Oversubscription**



## Use Unlimited Resources

The ASA lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available. For example, Context A, B, and C are in the Silver Class, which limits each class member to 1 percent of the connections, for a total of 3 percent; but the three contexts are currently only using 2 percent combined. Gold Class has unlimited access to connections. The contexts in the Gold Class can use more than the 97 percent of “unassigned” connections; they can also use the 1 percent of connections not currently in use by Context A, B, and C, even if that means that Context A, B, and C are unable to reach their 3 percent combined limit. (See [Figure 8-7](#).) Setting unlimited access is similar to oversubscribing the ASA, except that you have less control over how much you oversubscribe the system.

**Figure 8-7** Unlimited Resources



153211

## 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 [Assign MAC Addresses to Context Interfaces Automatically](#), page 8-21.

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 [How the ASA Classifies Packets](#), page 8-3 for information about classifying packets.

In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context. See [Configure the MAC Address, MTU, and TCP MSS](#), page 16-5 to manually set the MAC address.

- [Default MAC Address](#), page 8-12
- [Interaction with Manual MAC Addresses](#), page 8-12
- [Failover MAC Addresses](#), page 8-12
- [MAC Address Format](#), page 8-12

## Default MAC Address

(8.5(1.7) and Later) Automatic MAC address generation is enabled by default. The ASA autogenerates the prefix based on the last two bytes of the interface (ASA 5500-X) 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-X 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 [MAC Address Format, page 8-12](#).



### Note

(8.5(1.6) and earlier) To maintain hitless upgrade for failover pairs, the ASA does not convert an existing legacy auto-generation configuration upon a reload if failover is enabled. However, we strongly recommend that you manually change to the prefix method of generation when using failover, especially for the ASASM. Without the prefix method, ASASMs installed in different slot numbers experience a MAC address change upon failover, and can experience traffic interruption. After upgrading, to use the prefix method of MAC address generation, reenable MAC address autogeneration to use a prefix. For more information about the legacy method, see the **mac-address auto** command in the command reference.

## Interaction with Manual MAC Addresses

If you manually assign a MAC address and also enable auto-generation, then the manually assigned MAC address is used. If you later remove the manual MAC address, the auto-generated address is used.

Because auto-generated addresses (when using a prefix) start with A2, you cannot start manual MAC addresses with A2 if you also want to use auto-generation.

## Failover MAC Addresses

For use with failover, the ASA generates both an active and standby MAC address for each interface. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption. See [MAC Address Format, page 8-12](#) 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-X) 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 for Multiple Context Mode

| Model                                | License Requirement                                                                                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASA 5506-X                           | No support.                                                                                                                                                  |
| ASA 5508-X                           | Security Plus License: 2 contexts.<br><i>Optional license: 5 contexts.</i>                                                                                   |
| ASA 5512-X                           | <ul style="list-style-type: none"> <li>Base License: No support.</li> <li>Security Plus License: 2 contexts.</li> </ul> <i>Optional license: 5 contexts.</i> |
| ASA 5515-X                           | Base License: 2 contexts.<br><i>Optional license: 5 contexts.</i>                                                                                            |
| ASA 5516-X                           | Security Plus License: 2 contexts.<br><i>Optional license: 5 contexts.</i>                                                                                   |
| ASA 5525-X                           | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, or 20 contexts.</i>                                                                                |
| ASA 5545-X                           | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, 20, or 50 contexts.</i>                                                                            |
| ASA 5555-X                           | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, 20, 50, or 100 contexts.</i>                                                                       |
| ASA 5585-X with SSP-10               | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, 20, 50, or 100 contexts.</i>                                                                       |
| ASA 5585-X with SSP-20, -40, and -60 | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.</i>                                                                  |
| ASASM                                | Base License: 2 contexts.<br><i>Optional licenses: 5, 10, 20, 50, 100, or 250 contexts.</i>                                                                  |
| ASAv                                 | No support.                                                                                                                                                  |

## Prerequisites

After you are in multiple context mode, connect to 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 2, “Getting Started,”](#) for more information about connecting to the ASA.

## Guidelines for Multiple Context Mode

### 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.

---

### 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>).

## Defaults for Multiple Context Mode

- By default, the ASA is in single context mode.
- See [Default Class, page 8-9](#).

- See [Default MAC Address, page 8-12](#).

## Configure Multiple Contexts

To configure multiple context mode, perform the following steps:

- 
- |               |                                                                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable multiple context mode. See <a href="#">Enable or Disable Multiple Context Mode, page 8-15</a> .                                                                                                                                                     |
| <b>Step 2</b> | (Optional) Configure classes for resource management. See <a href="#">Configure a Class for Resource Management, page 8-16</a> . <b>Note:</b> For VPN support, you must configure VPN resources in a resource class; the default class does not allow VPN. |
| <b>Step 3</b> | Configure interfaces in the system execution space. <ul style="list-style-type: none"><li>• ASA 5500-X—<a href="#">Chapter 11, “Basic Interface Configuration.”</a></li><li>• ASASM—ASASM quick start guide.</li></ul>                                     |
| <b>Step 4</b> | Configure security contexts. See <a href="#">Configure a Security Context, page 8-18</a> .                                                                                                                                                                 |
| <b>Step 5</b> | (Optional) Customize MAC address assignments. See <a href="#">Assign MAC Addresses to Context Interfaces Automatically, page 8-21</a> .                                                                                                                    |
| <b>Step 6</b> | Complete interface configuration in the context. See <a href="#">Chapter 15, “Routed and Transparent Mode Interfaces.”</a>                                                                                                                                 |
- 

## Enable or Disable Multiple Context Mode

Your ASA might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you need to convert from single mode to multiple mode, follow the procedures in this section.

ASDM supports changing modes from single to multiple mode if you use the High Availability and Scalability Wizard and you enable Active/Active failover. See [Chapter 9, “Failover for High Availability,”](#) for more information. If you do not want to use Active/Active failover or want to change back to single mode, you must change modes using the CLI; because changing modes requires confirmation, you cannot use the Command Line Interface tool. This section describes changing modes at the CLI.

- [Enable Multiple Context Mode, page 8-15](#)
- [Restore Single Context Mode, page 8-16](#)

### Enable 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 [Manage Files, page 35-8](#).

### Procedure

- 
- Step 1** Change to multiple context mode:

```
mode multiple
```

Example:

```
ciscoasa(config)# mode multiple
```

You are prompted to reboot the ASA.

---

## Restore 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.

### Before You Begin

Perform this procedure in the system execution space.

### Procedure

- 
- Step 1** Copy the backup version of your original running configuration to the current startup configuration:

```
copy disk0:old_running.cfg startup-config
```

Example:

```
ciscoasa(config)# copy disk0:old_running.cfg startup-config
```

- Step 2** Set the mode to single mode:

```
mode single
```

Example:

```
ciscoasa(config)# mode single
```

You are prompted to reboot the ASA.

---

## Configure 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.

**Before You Begin**

- Perform this procedure in the system execution space.
- [Table 8-1](#) lists the resource types and the limits.

**Table 8-1 Resource Names and Limits**

| Resource Name                         | Rate or Concurrent    | Minimum and Maximum Number per Context | System Limit <sup>1</sup>                                                                                                                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------|-----------------------|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASDM Sessions                         | Concurrent            | 1 minimum<br>5 maximum                 | 32                                                                                                                                                                 | ASDM management sessions.<br><br><b>Note</b> ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 32 ASDM sessions represents a limit of 64 HTTPS sessions.                                                                                                                                                                                                           |
| Connections<br>Conns/sec <sup>2</sup> | Concurrent<br>or Rate | N/A                                    | Concurrent connections:<br>See <a href="#">Supported Feature Licenses Per Model</a> , page 3-1 for the connection limit available for your model.<br><br>Rate: N/A | TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts.                                                                                                                                                                                                                                                                                                                                                                                                             |
| Hosts                                 | Concurrent            | N/A                                    | N/A                                                                                                                                                                | Hosts that can connect through the ASA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Inspects/sec                          | Rate                  | N/A                                    | N/A                                                                                                                                                                | Application inspections per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MAC Entries                           | Concurrent            | N/A                                    | 65,535                                                                                                                                                             | For transparent firewall mode, the number of MAC addresses allowed in the MAC address table.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Routes                                | Concurrent            | N/A                                    | N/A                                                                                                                                                                | Dynamic routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Site-to-Site<br>VPN Burst             | Concurrent            | N/A                                    | The Other VPN session amount for your model minus the sum of the sessions assigned to all contexts for Site-to-Site VPN.                                           | The number of site-to-site VPN sessions allowed beyond the amount assigned to a context with Site-to-Site VPN. For example, if your model supports 5000 sessions, and you assign 4000 sessions across all contexts with Site-to-Site VPN, then the remaining 1000 sessions are available for Site-to-Site VPN Burst. Unlike Site-to-Site VPN, which guarantees the sessions to the context, Site-to-Site VPN Burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis. |
| Site-to-Site<br>VPN                   | Concurrent            | N/A                                    | See <a href="#">Supported Feature Licenses Per Model</a> , page 3-1 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.                                                                                                                                                                                                                                                                                                                |

**Table 8-1** *Resource Names and Limits (continued)*

| Resource Name       | Rate or Concurrent | Minimum and Maximum Number per Context | System Limit <sup>1</sup> | Description                   |
|---------------------|--------------------|----------------------------------------|---------------------------|-------------------------------|
| SSH                 | Concurrent         | 1 minimum<br>5 maximum                 | 100                       | SSH sessions.                 |
| Syslogs/sec         | Rate               | N/A                                    | N/A                       | Syslog messages per second.   |
| Telnet              | Concurrent         | 1 minimum<br>5 maximum                 | 100                       | Telnet sessions.              |
| xlates <sup>2</sup> | 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'").

### Procedure

- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Context Management > Resource Class**, and click **Add**.  
The **Add Resource Class** dialog box appears.
- Step 3** Enter a class name up to 20 characters in length, in the **Resource Class** field,.
- Step 4** In the **Count Limited Resources** area, set the concurrent limits for resources,.  
See [Table 8-1 on page 8-17](#) for a description of each resource type.  
For resources that do not have a system limit, you cannot set the percentage; you can only set an absolute value. If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then the resource is unlimited, or the system limit if available. For most resources, 0 sets the limit to unlimited. For VPN types, 0 sets the limit none.
- Step 5** In the **Rate Limited Resources** area, set the rate limit for resources.  
See [Table 8-1 on page 8-17](#) for a description of each resource type.  
If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then it is unlimited by default. 0 sets the limit to unlimited.
- Step 6** Click **OK**.

## Configure 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.

### Before You Begin

- Perform this procedure in the system execution space.

- For the ASASM, assign VLANs to the ASASM on the switch according to the ASASM quick start guide.
- For the ASA 5500-X, configure physical interface parameters, VLAN subinterfaces, EtherChannels, and redundant interfaces according to [Chapter 11, “Basic Interface Configuration.”](#)

### Procedure

- 
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Context Management > Security Contexts**, and click **Add**.  
The **Add Context** dialog box appears.
- Step 3** In the **Security Context** field, enter the context name as a string up to 32 characters long.  
This name is case sensitive, so you can have two contexts named “customerA” and “CustomerA,” for example. “System” or “Null” (in upper or lower case letters) are reserved names, and cannot be used.
- Step 4** In the **Interface Allocation** area, click the **Add** button to assign an interface to the context.
- From the **Interfaces > Physical** Interface drop-down list, choose an interface.  
You can assign the main interface, in which case you leave the subinterface ID blank, or you can assign a subinterface or a range of subinterfaces associated with this interface. In transparent firewall mode, only interfaces that have not been allocated to other contexts are shown. If the main interface was already assigned to another context, then you must choose a subinterface.
  - (Optional) In the **Interfaces > Subinterface Range** (optional) drop-down list, choose a subinterface ID.  
For a range of subinterface IDs, choose the ending ID in the second drop-down list, if available.  
In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.
  - (Optional) In the **Aliased Names** area, check **Use Aliased Name in Context** to set an aliased name for this interface to be used in the context configuration instead of the interface ID.
    - In the **Name** field, sets the aliased name.  
An aliased name must start with a letter, end with a letter, and have as interior characters only letters, digits, or an underscore. This field lets you specify a name that ends with a letter or underscore; to add an optional digit after the name, set the digit in the Range field.
    - (Optional) In the **Range** field, set the numeric suffix for the aliased name.  
If you have a range of subinterfaces, you can enter a range of digits to be appended to the name.
  - (Optional) To enable context users to see physical interface properties even if you set an aliased name, check **Show Hardware Properties in Context**.
  - Click **OK** to return to the Add Context dialog box.
- Step 5** (Optional) If you use IPS virtual sensors, then assign a sensor to the context in the IPS Sensor Allocation area.  
For detailed information about IPS and virtual sensors, see the firewall configuration guide.
- Step 6** (Optional) To assign this context to a resource class, choose a class name from the **Resource Assignment > Resource Class** drop-down list.  
You can add or edit a resource class directly from this area. See [Configure a Class for Resource Management, page 8-16](#) for more information.

- Step 7** Identify the URL by choosing a file system type from the **Config URL** drop-down list and entering a path in the field, to set the context configuration location.
- For example, the combined URL for FTP has the following format:
- ```
ftp://server.example.com/configs/admin.cfg
```
- a. (Optional) For external file systems, set the username and password by clicking **Login**.
- Step 8** (Optional) To set the failover group for Active/Active failover, choose the group name in the **Failover Group** drop-down list.
- Step 9** (Optional) To enable ScanSafe inspection in this context, click **Enable**. To override the license set in the system configuration, enter a license in the **License** field.
- Step 10** (Optional) In the **Description** field, add a description.
- Step 11** Click **OK** to return to the **Security Contexts** pane.

Configuration > Context Management > Security Contexts

Create, edit or delete security contexts.

Context	Mode	Interfaces	Primary...	Seconda...	Resou...	Config...	Group	Description
admin	Routed	Management0/0 Port-channel33			default	disk0:/...		
c10	Routed				default	disk0:/...		
c2	Routed	GigabitEthernet0/1.2-6 Management0/0			default	disk0:/...		
c3	Routed	GigabitEthernet0/2.1 GigabitEthernet0/2.3 GigabitEthernet0/2.5 Management0/0			default	disk0:/...		
c4	Routed	GigabitEthernet0/2.2 GigabitEthernet0/2.4 GigabitEthernet0/2.6 Management0/0			default	disk0:/...		
c5	Routed				default	disk0:/...		
c6	Routed				default	disk0:/...		
c7	Routed				default	disk0:/...		

☐ Enable auto-generation of MAC addresses for context interfaces that share a system interface


☐ Specify Pref...

Maximum TLS Sessions

☐ Specify the maximum number of TLS Proxy sessions that the ASA needs to support. By default, ASA supports 300 sessions.

Maximum Number of Sessions:

- Step 12** (Optional) To set the firewall mode to transparent, select the context and click **Change Firewall Mode**. You see the following confirmation dialog box:



If this is a new context, there is no configuration to erase. Click **Change Mode** to change to transparent firewall mode.

If this is an existing context, then be sure to back up the configuration before you change the mode.



Note You cannot change the mode of the currently connected context in ASDM (typically the admin context); see [Set the Firewall Mode \(Single Mode\)](#), page 6-9 to set the mode at the command line.

- Step 13** To customize auto-generation of MAC addresses, see [Assign MAC Addresses to Context Interfaces Automatically](#), page 8-21.
- Step 14** Check the **Specify the maximum number of TLS Proxy sessions that the ASA needs to support** check box, to specify the maximum TLS Proxy sessions for the device. For more information about TLS proxy, see the firewall configuration guide.

Assign MAC Addresses to Context Interfaces Automatically

This section describes how to configure auto-generation of MAC addresses.

The MAC address is used to classify packets within a context. See [About MAC Addresses](#), page 8-11 for more information, especially if you are upgrading from an earlier ASA version. See also [View Assigned MAC Addresses](#), page 8-27.

Before You Begin

- When you configure a name for the interface in a context, the new MAC address is generated immediately. If you enable this feature after you configure context interfaces, then MAC addresses are generated for all interfaces immediately after you enable it. If you disable this feature, the MAC address for each interface reverts to the default MAC address. For example, subinterfaces of GigabitEthernet 0/1 revert to using the MAC address of GigabitEthernet 0/1.
- In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context. See [Configure the MAC Address, MTU, and TCP MSS](#), page 16-5 to manually set the MAC address.

Procedure

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Context Management > Security Contexts**, and check **Mac-Address auto**. If you do not enter a prefix, then the ASA autogenerates the prefix based on the last two bytes of the interface (ASA 5500-X) or backplane (ASASM) MAC address.
- Step 3** (Optional) Check the **Prefix** check box, and in the field, enter a decimal value between 0 and 65535. This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address. See [MAC Address Format, page 8-12](#) section for more information about how the prefix is used.
-

Change Between Contexts and the System Execution Space

If you log in to the system execution space (or the admin context), you can change between contexts and perform configuration and monitoring tasks within each context. The running configuration that you edit in a configuration mode depends on your location. When you are in the system execution space, the running configuration consists only of the system configuration; when you are in a context, the running configuration consists only of that context.

Procedure

-
- Step 1** In the Device List pane, double-click **System** under the active device IP address, to configure the System.
- Step 2** In the Device List pane, double-click the context name under the active device IP address, to configure a context.
-

Manage Security Contexts

This section describes how to manage security contexts.

- [Remove a Security Context, page 8-22](#)
- [Change the Admin Context, page 8-23](#)
- [Change the Security Context URL, page 8-24](#)
- [Reload a Security Context, page 8-25](#)

Remove a Security Context

You cannot remove the current admin context.

**Note**

If you use failover, there is a delay between when you remove the context on the active unit and when the context is removed on the standby unit.

Before You Begin

Perform this procedure in the system execution space.

Procedure

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Context Management > Security Contexts**.
- Step 3** Select the context you want to delete, and click **Delete**.
The Delete Context dialog box appears.
- Step 4** If you might want to re-add this context later, and want to keep the configuration file for future use, uncheck the **Also delete config URL file from the disk** check box.
If you want to delete the configuration file, then leave the check box checked.
- Step 5** Click **Yes**.
-

Change the Admin Context

The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users.

**Note**

For ASDM, you cannot change the admin context within ASDM because your ASDM session would disconnect. You can perform this procedure using the Command Line Interface tool noting that you will have to reconnect to the new admin context.

Before You Begin

- You can set any context to be the admin context, as long as the configuration file is stored in the internal flash memory.
- Perform this procedure in the system execution space.

Procedure

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Tools > Command Line Interface**.
The Command Line Interface dialog box appears.
- Step 3** Enter the following command:

```
admin-context context_name
```

Step 4 Click **Send**.

Any remote management sessions, such as Telnet, SSH, or HTTPS (ASDM), that are connected to the admin context are terminated. You must reconnect to the new admin context.

**Note**

A few system configuration commands, including **ntp server**, identify an interface name that belongs to the admin context. If you change the admin context, and that interface name does not exist in the new admin context, be sure to update any system commands that refer to the interface.

Change the Security Context URL

This section describes how to change the context URL.

Before You Begin

- 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.
- Perform this procedure in the system execution space.

Procedure

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Context Management > Security Contexts**.
- Step 3** Select the context you want to edit, and click **Edit**.
The Edit Context dialog box appears.
- Step 4** Enter a new URL in the Config URL field, and click **OK**.
The system immediately loads the context so that it is running.
-

Reload 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.
- [Reload by Clearing the Configuration, page 8-25](#)
- [Reload by Removing and Re-adding the Context, page 8-25](#)

Reload by Clearing the Configuration

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

Procedure

-
Step 1	In the Device List pane, double-click the context name under the active device IP address.
Step 2	Choose Tools > Command Line Interface . The Command Line Interface dialog box appears.
Step 3	Enter the following command: <code>clear configure all</code>
Step 4	Click Send . The context configuration is cleared.
Step 5	Choose Tools > Command Line Interface again. The Command Line Interface dialog box appears.
Step 6	Enter the following command: <code>copy startup-config running-config</code>
Step 7	Click Send . The ASA reloads the configuration. The ASA copies the configuration from the URL specified in the system configuration. You cannot change the URL from within a context.
-

Reload 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. [Remove a Security Context, page 8-22](#). Be sure to uncheck the **Also delete config URL file from the disk** check box.

2. Configure a Security Context, page 8-18

Monitoring Security Contexts

This section describes how to view and monitor context information.

- [Monitor Context Resource Usage, page 8-26](#)
- [View Assigned MAC Addresses, page 8-27](#)

Monitor Context Resource Usage

To monitor resource usage of all contexts from the system execution space, perform the following steps:

-
- Step 1** If you are not already in the System mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Click the **Monitoring** button on the toolbar.
- Step 3** Click **Context Resource Usage**.

Click each resource type to view the resource usage for all contexts:

- **ASDM/Telnet/SSH**—Shows the usage of ASDM, Telnet, and SSH connections.
 - Context—Shows the name of each context.

For each access method, see the following usage statistics:

- Existing Connections (#)—Shows the number of existing connections.
 - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
 - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Routes**—Shows the usage of dynamic routes.
 - Context—Shows the name of each context.
 - Existing Connections (#)—Shows the number of existing connections.
 - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
 - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
 - **Xlates**—Shows the usage of network address translations.
 - Context—Shows the name of each context.
 - Xlates (#)—Shows the number of current xlates.
 - Xlates (%)—Shows the xlates used by this context as a percentage of the total number of xlates used by all contexts.
 - Peak (#)—Shows the peak number of xlates since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
 - **NATs**—Shows the number of NAT rules.

- Context—Shows the name of each context.
- NATs (#)—Shows the current number of NAT rules.
- NATs (%)—Shows the NAT rules used by this context as a percentage of the total number of NAT rules used by all contexts.
- Peak NATs (#)—Shows the peak number of NAT rules since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Syslogs**—Shows the rate of system log messages.
 - Context—Shows the name of each context.
 - Syslog Rate (#/sec)—Shows the current rate of system log messages.
 - Syslog Rate (%)—Shows the system log messages generated by this context as a percentage of the total number of system log messages generated by all contexts.
 - Peak Syslog Rate (#/sec)—Shows the peak rate of system log messages since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **VPN**—Shows the usage of VPN site-to-site tunnels.
 - Context—Shows the name of each context.
 - VPN Connections—Shows usage of guaranteed VPN sessions.
 - VPN Burst Connections—Shows usage of burst VPN sessions.
 - Existing (#)—Shows the number of existing tunnels.
 - Peak (#)—Shows the peak number of tunnels since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

Step 4 Click **Refresh** to refresh the view.

View Assigned MAC Addresses

You can view auto-generated MAC addresses within the system configuration or within the context.

- [View MAC Addresses in the System Configuration, page 8-27](#)
- [View MAC Addresses Within a Context, page 8-28](#)

View MAC Addresses in the System Configuration

This section describes how to view MAC addresses in the system configuration.

Before You Begin

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.

Procedure

- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

- Step 2** Choose **Configuration > Context Management > Security Contexts**, and view the Primary MAC and Secondary MAC columns.

View MAC Addresses Within a Context

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

Procedure

- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Choose **Configuration > Interfaces**, and view the MAC Address address column.
- This table shows the MAC address in use; if you manually assign a MAC address and also have auto-generation enabled, then you can only view the unused auto-generated address from within the system configuration.

History for Multiple Context Mode

Table 8-2 History for Multiple Context Mode

Feature Name	Platform Releases	Feature Information
Multiple security contexts	7.0(1)	Multiple context mode was introduced. We introduced the following screens: Configuration > Context Management.
Automatic MAC address assignment	7.2(1)	Automatic assignment of MAC address to context interfaces was introduced. We modified the following screen: Configuration > Context Management > Security Contexts.
Resource management	7.2(1)	Resource management was introduced. We introduced the following screen: Configuration > Context Management > Resource Management.
Virtual sensors for IPS	8.0(2)	The AIP SSM running IPS software Version 6.0 and above can run multiple virtual sensors, which means you can configure multiple security policies on the AIP SSM. You can assign each context or single mode ASA to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. We modified the following screen: Configuration > Context Management > Security Contexts.

Table 8-2 History for Multiple Context Mode (continued)

Feature Name	Platform Releases	Feature Information
Automatic MAC address assignment enhancements	8.0(5)/8.2(2)	The MAC address format was changed to use a prefix, to use a fixed starting value (A2), and to use a different scheme for the primary and secondary unit MAC addresses in a failover pair. The MAC addresses are also now persistent across reloads. The command parser now checks if auto-generation is enabled; if you want to also manually assign a MAC address, you cannot start the manual MAC address with A2. We modified the following screen: Configuration > Context Management > Security Contexts.
Maximum contexts increased for the ASA 5550 and 5580	8.4(1)	The maximum security contexts for the ASA 5550 was increased from 50 to 100. The maximum for the ASA 5580 was increased from 50 to 250.
Automatic MAC address assignment enabled by default	8.5(1)	Automatic MAC address assignment is now enabled by default. We modified the following screen: Configuration > Context Management > Security Contexts.
Automatic generation of a MAC address prefix	8.6(1)	In multiple context mode, the ASA now converts the automatic MAC address generation configuration to use a default prefix. The ASA auto-generates the prefix based on the last two bytes of the interface (ASA 5500-X) 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. If you want to change the prefix, you can reconfigure the feature with a custom prefix. The legacy method of MAC address generation is no longer available. Note To maintain hitless upgrade for failover pairs, the ASA does <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. We modified the following screen: Configuration > Context Management > Security Contexts
Dynamic routing in Security Contexts	9.0(1)	EIGRP and OSPFv2 dynamic routing protocols are now supported in multiple context mode. OSPFv3, RIP, and multicast routing are not supported.

Table 8-2 *History for Multiple Context Mode (continued)*

Feature Name	Platform Releases	Feature Information
New resource type for routing table entries	9.0(1)	A new resource type, routes, was created to set the maximum number of routing table entries in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class
Site-to-Site VPN in multiple context mode	9.0(1)	Site-to-site VPN tunnels are now supported in multiple context mode.
New resource type for site-to-site VPN tunnels	9.0(1)	New resource types, vpn other and vpn burst other, were created to set the maximum number of site-to-site VPN tunnels in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class



Failover for High Availability

This chapter describes how to configure Active/Standby or Active/Active failover to accomplish high availability of the Cisco ASA.

- [About Failover, page 7-1](#)
- [Licensing for Failover, page 7-24](#)
- [Guidelines for Failover, page 7-24](#)
- [Defaults for Failover, page 7-25](#)
- [Configure Active/Standby Failover, page 7-26](#)
- [Configure Active/Active Failover, page 7-26](#)
- [Configure Optional Failover Parameters, page 7-27](#)
- [Manage Failover, page 7-32](#)
- [Monitoring Failover, page 7-36](#)
- [History for Failover, page 7-39](#)

About Failover

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.

- [Failover Modes, page 7-2](#)
- [Failover System Requirements, page 7-2](#)
- [Failover and Stateful Failover Links, page 7-3](#)
- [MAC Addresses and IP Addresses in Failover, page 7-8](#)
- [Intra- and Inter-Chassis Module Placement for the ASA Services Module, page 7-9](#)
- [Stateless and Stateful Failover, page 7-12](#)
- [Transparent Firewall Mode Requirements for Failover, page 7-14](#)
- [Failover Health Monitoring, page 7-16](#)
- [Failover Times, page 7-18](#)
- [Configuration Synchronization, page 7-18](#)
- [About Active/Standby Failover, page 7-20](#)

- [About Active/Active Failover, page 7-21](#)

Failover Modes

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.
- 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.”

Related Topics

[Upgrade a Failover Pair or ASA Cluster, page 33-4](#)

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.

Related Topics

[Failover or ASA Cluster Licenses, page 3-22](#)

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 7-4](#)
- [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). The ASA does not support sharing interfaces between user data and the failover link even if different subinterfaces are configured for user data and failover. A separate physical, EtherChannel, or redundant interface must be used for the failover link.

For a redundant interface used as the failover link, see the following benefits for added redundancy:

- When a failover unit boots up, it alternates between the member interfaces to detect an active unit.
- If a failover unit stops receiving keepalive messages from its peer on one of the member interfaces, it switches to the other member interface.

For an EtherChannel used as the failover 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.

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 7-5
- [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. For an EtherChannel used as the 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.

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 the following figures 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 the following figures.

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 the following figure.

Figure 7-5 Connecting with a Secure Switch


236373

Scenario 4—Recommended

The most reliable failover configurations use a redundant interface on the failover link, as shown in the following figures.

Figure 7-6 Connecting with Redundant Interfaces

236375

Figure 7-7 Connecting with Inter-switch Links

236376

MAC Addresses and IP Addresses in Failover

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.

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. In single context mode, you can manually configure virtual MAC addresses.

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.

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.

Related Topics

- [About MAC Addresses, page 8-11](#)
- [Configure Active/Active Failover, page 7-26](#)

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.

- [Intra-Chassis Failover, page 7-9](#)
- [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.

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.

The following figure 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.

The following figure 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 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 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.
- TCP state bypass connections
- DHCP server address leases
- Multicast routing
- State information for modules, such as the ASA FirePOWER module.
- 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 for Failover

There are special considerations for failover when using transparent firewall mode.

- [Transparent Mode Requirements for Appliances, ASAv, page 7-15](#)
- [Transparent Mode Requirements for the ASA Services Module, page 7-15](#)

Transparent Mode Requirements for Appliances, ASAv

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 bpu
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 the ASA Services Module

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 with hello messages. When a unit does not receive three consecutive hello messages on the failover link, the unit sends LANTEST 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 2 polling periods, it runs interface tests. If all interface tests fail for an interface, but this same 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.

If you have a services module, such as the ASA FirePOWER SSP, then the ASA also monitors the health of the module over the backplane interface. Failure of the module is considered a unit failure and will trigger failover. This setting is configurable.

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).



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.

Interface Tests

The ASA uses the following interface tests:

1. **Link Up/Down test**—A test of the interface status. If the Link Up/Down test indicates that the interface is down, then the ASA considers it failed. If the status is Up, then the ASA performs the Network Activity test.
2. **Network Activity test**—A received network activity test. The purpose of this test is to generate network traffic using LANTEST messages to determine which (if either) unit has failed. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any packets during the test (up to 5 seconds), then the interface is considered operational. If one unit receives traffic and the other unit does not, then the unit that received no traffic is considered failed. If neither unit received traffic, then the ASA starts the ARP test.
3. **ARP test**—A reading of the unit ARP cache for the 10 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 ASA starts the ping test.
4. **Broadcast Ping test**—A ping test that consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the testing starts over again with the ARP test.

Interface Status

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.

Failover Times

The following table 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 various types of configuration synchronization.

- [Running Configuration Replication, page 7-18](#)
- [File Replication, page 7-19](#)
- [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.

**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.

File Replication

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

The ASA uses a cached file for the AnyConnect client profile stored in `cache:/stc/profiles`, and not the file stored in the flash file system. To replicate the AnyConnect client profile to the standby unit, perform one of the following:

- Enter the **write standby** command on the active unit.
- Reapply the profile on the active unit.
- Reload the standby unit.
- Local Certificate Authorities (CAs)
- ASA images
- ASDM images

Command Replication

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

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

Failure to enter the changes on the appropriate unit for command replication to occur causes the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

The following commands are replicated to the standby ASA:

- All configuration commands except for **mode**, **firewall**, and **failover lan unit**
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**

- **write memory**

The following commands are *not* replicated to the standby ASA:

- All forms of the **copy** command except for **copy running-config startup-config**
- All forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **show**
- **terminal pager** and **pager**

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 7-21](#)

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.

The following table 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 Events**

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.

About Active/Active Failover

This section describes Active/Active failover.

- [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.

The following table 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 Events**

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 for Failover

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. There are some exceptions to this rule. See the following table for precise licensing requirements for failover.

Model	License Requirement
ASA 5506-X series	- Active/Standby—Security Plus License. - Active/Active—No Support. Note Each unit must have the same encryption license.
ASA 5512-X through ASA 5555-X	- ASA 5512-X—Security Plus License. - Other models—Base License. Note Each unit must have the same encryption license; each unit must have the same 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.
ASAv	- Active/Standby—Standard and Premium Licenses. - Active/Active—No Support. Note The standby unit requires the same model license as the primary unit; Each unit must have the same encryption license.
All other models	Base License. Note Each unit must have the same encryption license.

Guidelines for Failover

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.

Model Support

For the ASA 5506W-X, you must disable interface monitoring for the internal GigabitEthernet 1/9 interface. These interfaces will not be able to communicate to perform the default interface monitoring checks, resulting in a switch from active to standby and back again because of expected interface communication failures.

Additional Guidelines

- 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 group 1 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.

Related Topics

[Auto Update Server Support in Failover Configurations, page 33-39](#)

Defaults for Failover

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 ASASM, all VLAN interfaces.

Configure Active/Standby Failover

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

Procedure

-
- Step 1** Choose **Wizards > High Availability and Scalability**. See select wizard guidelines in the following steps.
- Step 2** On the **Failover Peer Connectivity and Compatibility** screen, enter the IP address of the peer unit. This address must be an interface that has ASDM access enabled on it.
- By default, the peer address is assigned to be the standby address for the ASDM management interface.
- Step 3** On the **LAN Link Configuration** screen:
- **Active IP Address**—This IP address should be on an unused subnet.
 - **Standby IP Address**—This IP address must be on the same network as the active IP address.
 - (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note:** Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see [Modify the Failover Setup, page 7-32](#)).
- Step 4** On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:
- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link.
 - **Standby IP Address**—This IP address must be on the same network as the active IP address.
- Step 5** After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.
- After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.
- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
 - If you know failover is already enabled, click **Skip configuring peer**.
 - If you know the secondary unit is not yet failover-enabled, click **Stop waiting xx more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.
-

Configure Active/Active Failover

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

Procedure

-
- Step 1** Choose **Wizards > High Availability and Scalability**. See select wizard guidelines in the following steps.
- Step 2** In the **Failover Peer Connectivity and Compatibility Check** screen, the peer IP address must be an interface that has ASDM access enabled on it.

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

- Step 3** In the **Security Context Configuration** screen, if you converted to multiple context mode as part of the wizard, you will only see the admin context. You can add other contexts after you exit the wizard.
- Step 4** On the **LAN Link Configuration** screen:
- **Active IP Address**—This IP address should be on an unused subnet.
 - **Standby IP Address**—This IP address must be on the same network as the active IP address.
 - (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note:** Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see [Modify the Failover Setup, page 7-32](#)).
- Step 5** On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:
- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link.
 - **Standby IP Address**—This IP address must be on the same network as the active IP address.
- Step 6** After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.
- After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.
- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
 - If you know failover is already enabled, click **Skip configuring peer**.
 - If you know the secondary unit is not yet failover-enabled, click **Stop waiting xx more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.
-

Configure Optional Failover Parameters

You can customize failover settings as desired.

- [Configure Failover Criteria and Other Settings, page 7-27](#)
- [Configure Interface Monitoring and Standby Addresses, page 7-29](#)
- [Configure Support for Asymmetrically Routed Packets \(Active/Active Mode\), page 7-30](#)

Configure Failover Criteria and Other Settings

See [Defaults for Failover, 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. This section includes enabling HTTP replication per failover group for Active/Active mode; to configure HTTP replication for Active/Standby mode, see [Modify the Failover Setup, page 7-32](#).

Before You Begin

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

Procedure

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

Step 2 Disable the ability to make any configuration changes directly on the standby unit or context: on the **Setup** tab check the **Disable configuration changes on the standby unit** check box.

By default, configurations on the standby unit/context are allowed with a warning message.

Step 3 Click the **Criteria** tab.

Step 4 Configure the unit poll times:

In the **Failover Poll Times** area:

- **Unit Failover**—The amount of time between hello messages among units. The range is between 1 and 15 seconds or between 200 and 999 milliseconds.
- **Unit Hold Time**—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. The range is between 1 and 45 seconds or between 800 and 999 milliseconds. You cannot enter a value that is less than 3 times the polltime.



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

Step 5 (Active/Active mode only) Click the **Active/Active** tab, then choose a failover group and click **Edit**.

Step 6 (Active/Active mode only) Change the preferred role of the failover group: click either **Primary** or **Secondary**.

If you used the wizard, failover group 1 is assigned to the primary unit, and failover group 2 is assigned to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired.

Step 7 (Active/Active mode only) Configure failover group preemption: check the **Preempt after booting with optional delay of** check box.

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

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



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

Step 8 Configure the **Interface Policy**:

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



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

Step 9 (Active/Standby mode) Configure interface poll times:

In the **Failover Poll Time** area:

- **Monitored Interfaces**—The amount of time between polls among interfaces. The range is between 1 and 15 seconds or 500 to 999 milliseconds.
- **Interface Hold Time**—Sets the time during which a data interface must receive a hello message on the data interface, after which the peer is declared failed. Valid values are from 5 to 75 seconds.

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

Step 10 (Active/Active mode only) Enable HTTP replication: check the **Enable HTTP replication** check box.

For Active/Standby mode, see [Modify the Failover Setup, page 7-32](#). For both modes, see [Modify the Failover Setup, page 7-32](#) section for the HTTP replication rate.

Step 11 Configure virtual MAC addresses:

- Active/Standby mode—click the **MAC Addresses** tab, and click **Add**.
The **Add/Edit Interface MAC Address** dialog box appears.
- Active/Active mode—Go to the bottom of the **Active/Active** tab.

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

- Choose an interface from the **Physical Interface** drop-down list.
- In the **Active MAC Address** field, type the new MAC address for the active interface.
- In the **Standby MAC Address** field, type the new MAC address for the standby interface.
- Click **OK**. (Active/Active mode only) Click **OK** again.

Step 12 Click **Apply**.

Configure Interface Monitoring and Standby Addresses

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


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

Before You Begin

- 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.

Procedure

- Step 1** In single mode, choose **Configuration > Device Management > High Availability > Failover > Interfaces**.
- In multiple context mode, within a context choose **Configuration > Device Management > Failover > Interfaces**



A list of configured interfaces appears as well as any installed hardware modules, such as the ASA FirePOWER module. The **Monitored** column displays whether or not an interface is monitored as part of your failover criteria. If it is monitored, a check appears in the **Monitored** check box.

If you do not want a hardware module failure to trigger failover, you can disable module monitoring.

The IP address for each interface appears in the **Active IP Address** column. If configured, the standby IP address for the interface appears in the **Standby IP Address** column. The failover link and state link do not display IP address; you cannot change those addresses from this tab.

- Step 2** To disable monitoring of a listed interface, uncheck the **Monitored** check box for the interface.
- Step 3** To enable monitoring of a listed interface, check the **Monitored** check box for the interface.
- Step 4** For each interface that does not have a standby IP address, double-click the **Standby IP Address** field and enter an IP address into the field.
- Step 5** Click **Apply**.

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

When running in Active/Active failover, a unit might 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.

The following figure 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.

Before You Begin

- **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.
- You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group.

Procedure

-
- Step 1** On the primary unit active context, choose **Configuration > Device Setup > Routing > ASR Groups**.
 - Step 2** For the interface that receives asymmetrically routed packets, choose an **ASR Group ID** from the drop-down list.
 - Step 3** Click **Apply** to save your changes to the running configuration.
 - Step 4** Connect ASDM to the secondary unit, and choose the active context similar to the primary unit context.
 - Step 5** Choose **Configuration > Device Setup > Routing > ASR Groups**.
 - Step 6** For the similar interface on this unit, choose the same **ASR Group ID**.
 - Step 7** Click **Apply** to save your changes to the running configuration.
-

Manage Failover

This section describes how to manage failover units after you enable failover, including how to change the failover setup and how to force failover from one unit to another.

- [Modify the Failover Setup, page 7-32](#)
- [Force Failover, page 7-34](#)
- [Disable Failover, page 7-35](#)
- [Restore a Failed Unit, page 7-36](#)
- [Re-Sync the Configuration, page 7-36](#)

Modify the Failover Setup

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

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

Before You Begin

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

Procedure

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

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

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



Note Failover is not actually enabled until you apply your changes to the device.

Step 3 To encrypt communications on the failover and state links, use one of the following options:

- **IPsec Preshared Key** (preferred)—The preshared key is used by IKEv2 to establish IPsec LAN-to-LAN tunnels on the failover links between the failover units. Note: failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.
- **Secret Key**—Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

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

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

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

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

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

**Note**

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

If you choose an unconfigured interface or subinterface, you must supply the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** for the interface.

If you choose the failover link, you do not need to specify the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** values; the values specified for the failover link are used.

If you choose the **--Use Named--** option, the Logical Name field becomes a drop-down list of named interfaces. Choose the interface from this list. The **Active IP**, **Subnet Mask/Prefix Length**, and **Standby IP** values do not need to be specified. The values specified for the interface are used.

- **Logical Name**—Specify the logical name of the interface used for state communication, such as “state”. This name is informational.
- **Active IP**—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet, different from the failover link.
- **Standby IP**—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- **Subnet Mask**—Specify the subnet mask.
- (Optional, Active/Standby only) **Enable HTTP Replication**—This option enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover. In Active/Active mode, set the HTTP replication per failover group.

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

Step 7 Click **Apply**.

The configuration is saved to the device.

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

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

Related Topics

[Configure Failover Criteria and Other Settings, page 7-27](#)

Force Failover

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

Before You Begin

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

Procedure

-
- Step 1** To force failover at the unit level:
- a. Choose the screen depending on your context mode:
 - In single context mode choose **Monitoring > Properties > Failover > Status**.
 - In multiple context mode, in the System choose **Monitoring > Failover > System**.
 - b. Click one of the following buttons:
 - Click **Make Active** to make the unit this unit.
 - Click **Make Standby** to make the other unit the active unit.
- Step 2** (Active/Active mode only) To force failover at the failover group level:
- a. In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.
 - b. Click one of the following buttons:
 - Click **Make Active** to make the failover group active on this unit.
 - Click **Make Standby** to make the failover group active on the other unit.
-

Disable Failover

Disabling failover on one or both units causes the active and standby state of each unit to be maintained until you reload. For an Active/Active failover pair, the failover groups remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer.

See the following characteristics when you disable failover:

- The standby unit/context remains in standby mode so that both units do not start passing traffic (this is called a pseudo-standby state).
- The standby unit/context continues to use its standby IP addresses even though it is no longer connected to an active unit/context.
- The standby unit/context continues to listen for a connection on the failover link. If failover is re-enabled on the active unit/context, then the standby unit/context resumes ordinary standby status after re-synchronizing the rest of its configuration.
- To truly disable failover, save the no failover configuration to the startup configuration, and then reload.

Before You Begin

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

Procedure

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

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

- Step 2** Uncheck the **Enable Failover** check box.
 - Step 3** Click **Apply**.
 - Step 4** To completely disable failover, save the configuration and reload:
 - a. Click the **Save** button.
 - b. Choose **Tools > System Reload** and reload the ASA.
-

Restore a Failed Unit

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

Before You Begin

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

Procedure

-
- Step 1** To restore failover at the unit level:
 - a. Choose the screen depending on your context mode:
 - In single context mode choose **Monitoring > Properties > Failover > Status**.
 - In multiple context mode, in the System choose **Monitoring > Failover > System**.
 - b. Click **Reset Failover**.
 - Step 2** (Active/Active mode only) To reset failover at the failover group level:
 - a. In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.
 - b. Click **Reset Failover**.
-

Re-Sync the Configuration

Replicated commands are stored in the running configuration. To save replicated commands to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**.

Monitoring Failover

- [Failover Messages](#), page 7-37
- [Monitoring Failover Status](#), page 7-37

Failover Messages

When a failover occurs, both ASAs send out system messages.

- [Failover Syslog Messages, page 7-37](#)
- [Failover Debug Messages, page 7-37](#)
- [SNMP Failover Traps, page 7-37](#)

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 38, “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 37, “SNMP”](#) for more information.

Monitoring Failover Status

**Note**

After a failover event you should either re-launch ASDM or switch to another device in the Devices pane and then come back to the original ASA to continue monitoring the device. This action is necessary because the monitoring connection does not become re-established when ASDM is disconnected from and then reconnected to the device.

Choose **Monitoring > Properties > Failover** to monitor Active/Standby failover.

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

- [System, page 7-38](#)
- [Failover Group 1 and Failover Group 2, page 7-38](#)

System

The System pane displays the failover state of the system. You can also control the failover state of the system by:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

Fields

Failover state of the system—*Display only*. Displays the failover state of the ASA. The information shown is the same output you would receive from the **show failover** command. Refer to the command reference for more information about the displayed output.

The following actions are available on the System pane:

- **Make Active**—Click this button to make the ASA the active unit in an active/standby configuration. In an active/active configuration, clicking this button causes both failover groups to become active on the ASA.
- **Make Standby**—Click this button to make the ASA the standby unit in an active/standby pair. In an active/active configuration, clicking this button causes both failover groups to go to the standby state on the ASA.
- **Reset Failover**—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- **Reload Standby**—Click this button to force the standby unit to reload.
- **Refresh**—Click this button to refresh the status information in the Failover state of the system field.

Failover Group 1 and Failover Group 2

The Failover Group 1 and Failover Group 2 panes display the failover state of the selected group. You can also control the failover state of the group by toggling the active/standby state of the group or by resetting a failed group.

Fields

Failover state of Group[x]—*Display only*. Displays the failover state of the selected failover group. The information shown is the same as the output you would receive from the **show failover group** command.

You can perform the following actions from this pane:

- **Make Active**—Click this button to make the failover group active unit on the ASA.
- **Make Standby**—Click this button to force the failover group into the standby state on the ASA.
- **Reset Failover**—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- **Refresh**—Click this button to refresh the status information in the Failover state of the system field.

History for Failover

Table 7-4 *History for Failover*

Feature Name	Releases	Feature Information
Active/Standby failover	7.0(1)	This feature was introduced.
Active/Active failover	7.0(1)	This feature was introduced.
Support for a hex value for the failover key	7.0(4)	You can now specify a hex value for failover link encryption. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.
Support for the master passphrase for the failover key	8.3(1)	The failover key now supports the master passphrase, which encrypts the shared key in the running and startup configuration. If you are copying the shared secret from one ASA to another, for example from the more system:running-config command, you can successfully copy and paste the encrypted shared key. Note The failover key shared secret shows as ***** in show running-config output; this obscured key is not copyable. There were no ASDM changes.
IPv6 support for failover added.	8.2(2)	We modified the following screens: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	9.1(2)	Instead of using the proprietary encryption for the failover key, you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. Note Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.

Table 7-4 *History for Failover*

Feature Name	Releases	Feature Information
Disable health monitoring of a hardware module	9.3(1)	By default, the ASA monitors the health of an installed hardware module such as the ASA FirePOWER module. If you do not want a hardware module failure to trigger failover, you can disable module monitoring. We modified the following screen: Configuration > Device Management > High Availability and Scalability > Failover > Interfaces
Lock configuration changes on the standby unit or standby context in a failover pair	9.3(2)	You can now lock configuration changes on the standby unit (Active/Standby failover) or the standby context (Active/Active failover) so you cannot make changes on the standby unit outside normal configuration syncing. We modified the following screen: Configuration > Device Management > High Availability and Scalability > Failover > Setup



ASA Cluster

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 [Unsupported Features with Clustering](#), page 8-24.

- [About ASA Clustering](#), page 8-1
- [Licensing for ASA Clustering](#), page 8-30
- [Prerequisites for ASA Clustering](#), page 8-31
- [Guidelines for ASA Clustering](#), page 8-32
- [Defaults for ASA Clustering](#), page 8-35
- [Configure ASA Clustering](#), page 8-35
- [Manage ASA Cluster Members](#), page 8-50
- [Monitoring the ASA Cluster](#), page 8-58
- [Examples for ASA Clustering](#), page 8-60
- [History for ASA Clustering](#), page 8-73

About ASA Clustering

This section describes the clustering architecture and how it works.

- [How the ASA Cluster Fits into Your Network](#), page 8-2
- [Performance Scaling Factor](#), page 8-2
- [Cluster Members](#), page 8-3
- [Cluster Interfaces](#), page 8-4
- [Cluster Control Link](#), page 8-5
- [High Availability Within the ASA Cluster](#), page 8-8
- [Configuration Replication](#), page 8-10
- [ASA Cluster Management](#), page 8-11

- [Load Balancing Methods, page 8-12](#)
- [Inter-Site Clustering, page 8-17](#)
- [How the ASA Cluster Manages Connections, page 8-21](#)
- [ASA Features and Clustering, page 8-23](#)

How the ASA Cluster Fits into Your Network

The cluster consists of multiple ASAs acting as a single unit. To act as a cluster, the ASAs need the following infrastructure:

- Isolated, high-speed backplane network for intra-cluster communication, known as the *cluster control link*.
- Management access to each ASA for configuration and monitoring.

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.
- Policy-Based Routing (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using route maps and ACLs.
- 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.

Related Topics

- [Licensing for ASA Clustering, page 8-30](#)
- [Cluster Control Link, page 8-5](#)
- [ASA Cluster Management, page 8-11](#)
- [Spanned EtherChannel \(Recommended\), page 8-12](#)
- [Policy-Based Routing \(Routed Firewall Mode Only\), page 8-16](#)
- [Equal-Cost Multi-Path Routing \(Routed Firewall Mode Only\), page 8-17](#)

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

Cluster members work together to accomplish the sharing of the security policy and traffic flows. This section describes the nature of each member role.

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

Bootstrap Configuration

On each device, you configure a minimal bootstrap configuration including the cluster name, cluster control link interface, and other cluster settings. The first unit on which you enable clustering typically becomes the *master* unit. When you enable clustering on subsequent units, they join the cluster as *slaves*.

Master and Slave Unit Roles

One member of the cluster is the master unit. The master unit is determined by the priority setting in the bootstrap configuration; the priority is set between 1 and 100, where 1 is the highest priority. All other members are slave units. Typically, when you first create a cluster, the first unit you add becomes the master unit simply because it is the only unit in the cluster so far.

You must perform all configuration (aside from the bootstrap configuration) on the master unit only; the configuration is then replicated to the slave units. In the case of physical assets, such as interfaces, the configuration of the master unit is mirrored on all slave units. For example, if you configure GigabitEthernet 0/1 as the inside interface and GigabitEthernet 0/0 as the outside interface, then these interfaces are also used on the slave units as inside and outside interfaces.

Some features do not scale in a cluster, and the master unit handles all traffic for those features.

Related Topics

- [Centralized Features for Clustering, page 8-24](#)

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.

Related Topics


- [Centralized Features for Clustering, page 8-24](#)

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.

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.



333361


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.



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.



Related Topics

- [Load Balancing Methods, page 8-12](#)

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link.

- [Cluster Control Link Traffic Overview, page 8-6](#)

- [Cluster Control Link Interfaces and Network, page 8-6](#)
- [Size the Cluster Control Link, page 8-7](#)
- [Cluster Control Link Redundancy, page 8-7](#)
- [Cluster Control Link Reliability, page 8-8](#)
- [Cluster Control Link Failure, page 8-8](#)

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Master election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Related Topics

- [Cluster Members, page 8-3](#)
- [Configuration Replication, page 8-10](#)
- [Unit Health Monitoring, page 8-9](#)
- [Data Path Connection State Replication, page 8-10](#)
- [Rebalancing New TCP Connections Across the Cluster, page 8-23](#)

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.

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.

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.

Related Topics

- [Cluster Control Link Redundancy, page 8-7](#)
- [Size the Cluster Control Link, page 8-7](#)

Size 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.


Related Topics

- [Inter-Site Clustering, page 8-17.](#)

Cluster Control Link Redundancy

We recommend using an EtherChannel for the cluster control link, so that 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

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.



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.

Related Topics

[Rejoining the Cluster, page 8-10](#)

High Availability Within the ASA Cluster

ASA Clustering provides high availability by monitoring unit and interface health and by replicating connection states between units.

- [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.

When you enable health monitoring, all physical interfaces (including the main EtherChannel and redundant interface types) are monitored by default; you can optionally disable monitoring per interface.

Unit or Interface Failure

When health monitoring is enabled, a unit is removed from the cluster if it fails or if its monitored interfaces fail. If all physical ports for a given logical interface fail on a particular unit, but there are active ports under the same logical interface 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. The ASA does not monitor interfaces for the first 90 seconds that a unit joins the cluster. Interface status changes during this time will not cause the ASA to be removed from the cluster. 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.



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.

Related Topics

[Rejoining the Cluster, page 8-10](#)

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.
- **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.
- **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.

Related Topics

- [Configure ASA Cluster Parameters, page 8-50](#)

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 the following table 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

One of the benefits of using ASA clustering is the ease of management. This section describes how to manage the cluster.

- [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 that 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 that 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.

Related Topics

- [Chapter 18, “Digital Certificates.”](#)

Load Balancing Methods

Available load balancing methods depend on the firewall mode and interface type.

- [Spanned EtherChannel \(Recommended\), page 8-12](#)
- [Policy-Based Routing \(Routed Firewall Mode Only\), page 8-16](#)
- [Equal-Cost Multi-Path Routing \(Routed Firewall Mode Only\), page 8-17](#)

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 8-13](#)
- [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.

Related Topics

[EtherChannels, page 12-2](#)

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. On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS or Cisco 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.

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.

Related Topics

- [Customize the EtherChannel, page 12-10](#)
- [Load Balancing, page 12-4](#)
- [NAT and Clustering, page 8-28](#)

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.

Depending on your switches, you can configure up to 32 active links in the spanned EtherChannel. This feature requires both switches in the vPC to support EtherChannels with 16 active links each (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).


For switches that support 8 active links in the EtherChannel, you can configure up to 16 active links in the spanned EtherChannel when connecting to two switches in a VSS/vPC.

If you want to use more than 8 active links in a spanned EtherChannel, you cannot also have standby links; the support for 9 to 32 active links requires you to disable cLACP dynamic port priority that allows the use of standby links. You can still use 8 active links and 8 standby links if desired, for example, when connecting to a single switch.


The following figure shows a 32 active link spanned EtherChannel in an 8-ASA cluster and a 16-ASA cluster.



The following figure shows a 16 active link spanned EtherChannel in a 4-ASA cluster and an 8-ASA cluster.



The following figure shows a traditional 8 active/8 standby link spanned EtherChannel in a 4-ASA cluster and an 8-ASA cluster. 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 Cisco IOS PBR with Object Tracking. Cisco 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 URL for more details:

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml



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

For inter-site installations, you can take advantage of ASA clustering as long as you follow these guidelines.

- [Inter-Site Clustering Guidelines, page 8-17](#)
- [Sizing the Data Center Interconnect, page 8-18](#)
- [Inter-Site Examples, page 8-19](#)

Inter-Site Clustering Guidelines

See the following guidelines for inter-site clustering:

- Supports inter-site clustering in the following interface and firewall modes:

Interface Mode	Firewall Mode	
	Routed	Transparent
Individual Interface	Yes	N/A
Spanned EtherChannel	No	Yes

- 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; 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. This is expected behavior.

- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.
- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.

Related Topics

- [Rebalancing New TCP Connections Across the Cluster, page 8-23](#)
- [Connection Roles, page 8-22](#)

Sizing the Data Center Interconnect

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

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

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

For example:

- For 4 members at 2 sites:
 - 4 cluster members total
 - 2 members at each site
 - 5 Gbps cluster control link per member
 Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).
- For 8 members at 2 sites, the size increases:
 - 8 cluster members total
 - 4 members at each site
 - 5 Gbps cluster control link per member
 Reserved DCI bandwidth = 10 Gbps (4/2 x 5 Gbps).
- For 6 members at 3 sites:
 - 6 cluster members total
 - 3 members at site 1, 2 members at site 2, and 1 member at site 3
 - 10 Gbps cluster control link per member
 Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).

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

Reserved DCI bandwidth = 10 Gbps ($1/2 \times 10 \text{ Gbps} = 5 \text{ Gbps}$; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).


Inter-Site Examples

The following examples show supported cluster deployments.

- [Individual Interface Routed Mode North-South Inter-Site Example, page 8-19](#)
- [Spanned EtherChannel Transparent Mode North-South Inter-Site Example, page 8-20](#)
- [Spanned EtherChannel Transparent Mode East-West Inter-Site Example, page 8-21](#)

Individual Interface Routed Mode North-South Inter-Site Example


The following example shows 2 ASA cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). 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.



Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 ASA cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each ASA EtherChannel is spanned across all ASAs in the cluster.

The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. 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. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. 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.




The implementation of the switches at each site can include:

- **Inter-site VSS/vPC**—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the ASA cluster units at each Data Center to only connect to the local switch, while the VSS/vPC traffic goes across the DCI. In this case, connections are for the most part kept local to each datacenter. You can optionally connect each ASA unit to both switches across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.
- **Local VSS/vPC at each site**—For better switch redundancy, you can install 2 separate VSS/vPC pairs at each site. In this case, although the ASAs still have a spanned EtherChannel with Data Center 1 ASAs connected only to both local switches, and Data Center 2 ASAs connected to those local switches, the spanned EtherChannel is essentially “split.” Each local VSS/vPC sees the spanned EtherChannel as a site-local EtherChannel.

Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 ASA cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each ASA EtherChannel is spanned across all ASAs in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.



See [Spanned EtherChannel Transparent Mode North-South Inter-Site Example, page 8-20](#) for information about vPC/VSS options.

How the ASA Cluster Manages Connections

Connections can be load-balanced to multiple members of the cluster. Connection roles determine how to connections are handled in both normal operation and in a high availability situation.

- [Connection Roles, page 8-22](#)
- [New Connection Ownership, page 8-22](#)

- [Sample Data Flow, page 8-22](#)
- [Rebalancing New TCP Connections Across the Cluster, page 8-23](#)

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.

Related Topics

- [Load Balancing Methods, page 8-12](#)

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

Some ASA features are not supported with ASA clustering, and some are only supported on the master unit. Other features might have caveats for proper usage.

- [Unsupported Features with Clustering, page 8-24](#)
- [Centralized Features for Clustering, page 8-24](#)

- [Features Applied to Individual Units, page 8-25](#)
- [Dynamic Routing and Clustering, page 8-26](#)
- [Multicast Routing and Clustering, page 8-27](#)
- [NAT and Clustering, page 8-28](#)
- [AAA for Network Access and Clustering, page 8-29](#)
- [Syslog and NetFlow and Clustering, page 8-29](#)
- [SNMP and Clustering, page 8-29](#)
- [VPN and Clustering, page 8-29](#)
- [FTP and Clustering, page 8-30](#)
- [Cisco TrustSec and Clustering, page 8-30](#)

Unsupported Features with Clustering

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
 - SCCP (Skinny)
 - WAAS
 - WCCP
- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, and proxy. DHCP relay is supported.
- VPN load balancing
- Failover
- ASA CX module

Centralized Features for Clustering

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.

If you use the rebalancing feature, 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

Related Topics

- [Size the Cluster Control Link, page 8-7](#)
- [Rebalancing New TCP Connections Across the Cluster, page 8-23](#)

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.
- ASA FirePOWER module—There is no configuration sync or state sharing between ASA FirePOWER modules. You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the cluster using FireSIGHT Management Center. Do not use different ASA-interface-based zone definitions for devices in the cluster.
- ASA 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 and Clustering

- [Dynamic Routing in Spanned EtherChannel Mode, page 8-26](#)
- [Dynamic Routing in Individual Interface Mode, page 8-27](#)

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. See the OSPF Non-Stop Forwarding feature to address the interruption.

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 and Clustering

Multicast routing behaves differently depending on the interface mode.

- [Multicast Routing in Spanned EtherChannel Mode, page 8-27](#)
- [Multicast Routing in Individual Interface Mode, page 8-28](#)

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 and Clustering

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

SIP Inspection and Clustering

A control flow can be created on any unit (due to load balancing), but its child data flows must reside on the same unit.

TLS Proxy configuration is not supported.

AAA for Network Access and Clustering

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 and Clustering

- **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.
- **NetFlow**—Each unit in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

Related Topics

- [Include the Device ID in Non-EMBLEM Format Syslog Messages, page 38-21](#)

SNMP and Clustering

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 and Clustering

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 and Clustering

- 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 and Clustering

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 for ASA Clustering

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.

There are exceptions to this rule. See the following table for precise licensing requirements for clustering.

Model	License Requirement
ASA 5585-X	Cluster License, supports up to 16 units. Note Each unit must have the same encryption license; each unit must have the same 10 GE I/O/Security Plus license (ASA 5585-X with SSP-10 and -20).
ASA 5512-X	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	Base License, supports 2 units. Note Each unit must have the same encryption license.
All other models	No support.

Prerequisites for ASA Clustering

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.
- You can have cluster members in different geographical locations (inter-site) when using individual interface mode.
- 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.

Switch Prerequisites

- Be sure to complete the switch configuration before you configure clustering on the ASAs.
- For a list of supported switches, see [Cisco ASA Compatibility](#).

ASA Prerequisites

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

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.

Related Topics

- [Guidelines for ASA Clustering, page 8-32](#)
- [Enable Jumbo Frame Support, page 9-8](#)
- [Bootstrap Configuration, page 8-3](#)

Guidelines for ASA Clustering

Context Mode

The mode must match on each member unit.

Firewall Mode

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

Failover

Failover is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Models

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.

- ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X

Switches

- 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 Cisco Nexus OS and Cisco IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the ASAs in a cluster. *Do not* change the load-balancing algorithm from the default on the ASA .
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.
- Some switches do not support dynamic port priority with LACP (active and standby links). You can disable dynamic port priority to provide better compatibility with spanned EtherChannels.
- Network elements on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.


- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the ASA to fixed:

```
router(config)# port-channel id hash-distribution fixed
```


Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

EtherChannels

- 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.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
 - Spanned EtherChannels—For ASA *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.



- Device-local EtherChannels—For ASA *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple ASA EtherChannels into one EtherChannel on the switch.



Additional Guidelines

- 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.
- We do not support VXLAN in Individual Interface mode. Only Spanned EtherChannel mode supports VXLAN.

Related Topics

- [Size the Cluster Control Link, page 8-7](#)
- [Bootstrap Configuration, page 8-3](#)

- [Unsupported Features with Clustering](#), page 8-24
- [Configure an EtherChannel](#), page 12-8
- [Guidelines for EtherChannels and Redundant Interfaces](#), page 12-4

Defaults for ASA Clustering

- 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. Interface health monitoring is enabled on all interfaces by default.
- Connection rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.

Configure ASA Clustering

**Note**

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

To configure clustering, perform the following tasks:

-
- Step 1** Complete all pre-configuration on the switches and ASAs according to [Prerequisites for ASA Clustering](#), page 8-31 and [Guidelines for ASA Clustering](#), page 8-32.
 - Step 2** [Cable the Cluster Units and Configure Upstream and Downstream Equipment](#), page 8-35.
 - Step 3** [Back Up Your Configurations \(Recommended\)](#), page 8-37.
 - Step 4** [Configure the Cluster Interface Mode on the Master Unit](#), page 8-38. You can only configure one type of interface for clustering: Spanned EtherChannels or Individual interfaces.
 - Step 5** [\(Recommended; Required in Multiple Context Mode\) Configure Interfaces on the Master Unit](#), page 8-41. You cannot enable clustering if the interfaces are not cluster-ready. In single context mode, you can alternatively configure many interface settings within the High Availability and Scalability wizard, but not all interface options are available in the wizard, and you cannot configure the interfaces in multiple context mode within the wizard.
 - Step 6** [Create or Join an ASA Cluster](#), page 8-47.
 - 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.
-

Cable the Cluster Units and Configure 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.

Back Up Your Configurations (Recommended)

When you enable clustering on a slave unit, the current configuration is replaced with one synced from the master unit. If you ever want to leave the cluster entirely, it may be useful to have a backup configuration with a usable management interface configuration.

Before You Begin

Perform a backup on each unit.

Procedure

-
- Step 1** Choose **Tools > Backup Configurations**.
- Step 2** Back up at least the running configuration. See [Back Up the Local CA Server, page 35-22](#) for a detailed procedure.
-

Related Topics

- [Leave the Cluster, page 8-56](#)

Configure the Cluster Interface Mode on the Master Unit

You can only configure one type of interface for clustering: Spanned EtherChannels or Individual interfaces; you cannot mix interface types in a cluster.

**Note**

If you do not add slave units from the master unit, you must set the interface mode manually on all units according to this section, not just the master unit; if you add slaves from the master, ASDM sets the interface mode automatically on the slave.

Before You Begin

- 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.

Procedure**Step 1**

In ASDM on the master unit, choose **Tools > Command Line Interface**. Show any incompatible configuration so that you can force the interface mode and fix your configuration later; the mode is not changed with this command:

```
cluster interface-mode {individual | spanned} check-details
```

Example:

Command Line Interface

Type a command to be sent directly to the device. For command help, type a command followed by a question mark. For commands that would prompt for confirmation, add an appropriate noconfirm option as parameter to the command and send it to the device. To make the changes permanent, use the File > Save Running Configuration to Flash menu option to save the configuration to flash.

Command

☒ Single Line ☐ Multiple Line ☒ Enable context sensitive help (?)

cluster interface-mode spanned check-details

Response:

Result of the command: "cluster interface-mode spanned check-details"

ERROR: Please modify the following configuration elements that are incompatible with 'spanned' interface-mode.

- A cluster IP address pool must be specified on interface Gi0/0(outside). Or remove IP address configuration.
- A cluster IP address pool must be specified on interface Ma0/0(management). Or remove IP address configuration.

Clear Response

Help Close Send

**Caution**


After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool), you will not be able to reconnect because cluster-incompatible interface configuration is removed. In that case, you will have to connect to the console port to fix the interface configuration.

Step 2

Set the interface mode for clustering:

```
cluster interface-mode {individual | spanned} force
```

Example:




There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering.

The **force** option changes the mode without checking your configuration for incompatible settings. You need to manually fix any configuration issues after you change the mode. Because any interface configuration can only be fixed after you set the mode, we recommend using the **force** option so that you can at least start from the existing configuration. You can re-run the **check-details** option after you set the mode for more guidance.

Without the **force** option, if there is any incompatible configuration, you are prompted to clear your configuration and reload, thus requiring you to connect to the console port to reconfigure your management access. If your configuration is compatible (rare), the mode is changed and the configuration is preserved. If you do not want to clear your configuration, you can exit the command by typing **n**.

To remove the interface mode, enter the **no cluster interface-mode** command.

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



Related Topics

- [\(Recommended; Required in Multiple Context Mode\) Configure Interfaces on the Master Unit, page 8-41](#)

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

You must modify any interface that is currently configured with an IP address to be cluster-ready *before* you enable clustering. At a minimum, you must modify the management interface to which ASDM is currently connected. For other interfaces, you can configure them before or after you enable clustering; we recommend pre-configuring all of your interfaces so that the complete configuration is synced to new cluster members. In multiple context mode, you must use the procedures in this section to fix existing interfaces or to configure new interfaces. However, in single mode, you can skip this section and configure common interface parameters within the High Availability and Scalability wizard (see [Create or Join an ASA Cluster, page 8-47](#)). Note that advanced interface settings such as creating EtherChannels for Individual interfaces are not available in the wizard.

This section describes how to configure interfaces to be compatible with clustering. You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. Each method uses a different load-balancing mechanism. You cannot configure both types in the same configuration, with the exception of the management interface, which can be an Individual interface even in Spanned EtherChannel mode.

- [Configure Individual Interfaces \(Recommended for the Management Interface\), page 8-41](#)
- [Configure Spanned EtherChannels, page 8-44](#)

Related Topics

- [Cluster Interfaces, page 8-4](#)

Configure 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.

Before You Begin

- Except for the management-only interface, you must be in Individual interface mode.
- For multiple context mode, perform this procedure in each context. If you are not already in the context configuration mode in the Configuration > Device List pane, double-click the context name under the active device IP address.
- Individual interfaces require you to configure load balancing on neighbor devices. External load balancing is not required for the management interface.
- (Optional) Configure the interface as a device-local EtherChannel interface, a redundant interface, and/or configure subinterfaces.
 - For an EtherChannel, this EtherChannel is local to the unit, and is not a Spanned EtherChannel.

- Management-only interfaces cannot be redundant interfaces.
 - If you are connecting remotely to the management interface using ASDM, the current IP address of prospective slave units are for temporary use.
 - Each member will be assigned an IP address from the cluster IP pool defined on the master unit.
 - The cluster IP pool cannot include addresses already in use on the network, including prospective slave IP addresses.
- For example:
- a. You configure the master unit to use 10.1.1.1.
 - b. Other units use 10.1.1.2, 10.1.1.3, and 10.1.1.4.
 - c. When you configure the cluster IP pool on the master unit, you cannot include the .2, .3, or .4 addresses in the pool, because they are in use.
 - d. Instead, you need to use other IP addresses on the network, such as .5, .6, .7, and .8.




Note The pool needs as many addresses as there are members of the cluster, including the master unit; the original .1 address is the main cluster IP address that belongs to the current master unit.

- e. After you join the cluster, the old, temporary addresses are relinquished and can be used elsewhere.

Procedure

- Step 1** Choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- Step 2** Choose the interface row, and click **Edit**. Set the interface parameters. See the following guidelines:
 - (Required for a management interface in Spanned EtherChannel mode) **Dedicate this interface to management only**—Sets an interface to management-only mode so that it does not pass through traffic. By default, Management type interfaces are configured as management-only. In transparent mode, this command is always enabled for a Management type interface.
 - **Use Static IP**—DHCP and PPPoE are not supported.
- Step 3** To add the IPv4 cluster IP pool, and optionally a MAC address pool, click the **Advanced** tab.
 - a. In the **ASA Cluster** area, create a cluster IP pool by clicking the ... button next to the **IP Address Pool** field. The valid range shown is determined by the Main IP address you set on the General tab.
 - b. Click **Add**.
 - c. Configure a range of addresses that does not include the Main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.




- d. Click **OK** to create the new pool.
- e. Select the new pool you created, and click **Assign**, and then click **OK**.
The pool name appears in the **IP Address Pool** field.

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

- a. Check the **Enable IPv6** check box.
- b. In the **Interface IPv6 Addresses** area, click **Add**.
The **Enable address autoconfiguration** option is not supported.
The **Add IPv6 Address for Interface** dialog box appears.

- c. In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48. Click the ... button to configure the cluster IP pool.
- d. Click **Add**.



- e. Configure the starting IP address (network prefix), prefix length, and number of addresses in the pool.
- f. Click **OK** to create the new pool.
- g. Select the new pool you created, and click **Assign**, and then click **OK**.
The pool appears in the **ASA Cluster IP Pool** field.
- h. Click **OK**.

Step 5 Click **OK** to return to the Interfaces pane.

Step 6 Click **Apply**.

Related Topics

- [Management Interface, page 8-11](#)
- [Configure the Cluster Interface Mode on the Master Unit, page 8-38](#)
- [Load Balancing Methods, page 8-12](#)
- [Configure an EtherChannel, page 12-8](#)
- [Enable Jumbo Frame Support, page 9-8](#)
- [Configure VLAN Subinterfaces and 802.1Q Trunking, page 13-3](#)
-

Configure Spanned EtherChannels

A Spanned EtherChannel spans all ASAs in the cluster, and provides load balancing as part of the EtherChannel operation.

Before You Begin

- You must be in Spanned EtherChannel interface mode.
- For multiple context mode, start this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.
- For transparent mode, configure the bridge group.
- *Do not* specify the maximum and minimum links in the EtherChannel—We recommend that you do not specify the maximum and minimum links in the EtherChannel on either the ASA or the switch. If you need to use them, note the following:
 - The maximum links set on the ASA is the total number of active ports for the whole cluster. Be sure the maximum links value configured on the switch is not larger than the ASA value.
 - The minimum links set on the ASA is the minimum active ports to bring up a port-channel interface *per unit*. On the switch, the minimum links is the minimum links across the cluster, so this value will not match the ASA value.
- *Do not* change the load-balancing algorithm from the default. On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the ASAs in a cluster.
- When using Spanned EtherChannels, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

-
- Step 1** Depending on your context mode:
- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.
- Step 2** Choose **Add > EtherChannel Interface**.
- The **Add EtherChannel Interface** dialog box appears.

Step 3 Enable the following:

- **Port Channel ID**
- **Span EtherChannel across the ASA cluster**
- **Enable Interface** (checked by default)
- **Members in Group**—In the **Members in Group** list, you need to add at least one interface. Multiple interfaces in the EtherChannel per unit are useful for connecting to switches in a VSS or vPC. Keep in mind that by default, a spanned EtherChannel can have only 8 active interfaces out of 16 maximum across all members in the cluster; the remaining 8 interfaces are on standby in case of link failure. To use more than 8 active interfaces (but no standby interfaces), disable dynamic port priority. When you disable dynamic port priority, you can use up to 32 active links across the cluster. For example, for a cluster of 16 ASAs, you can use a maximum of 2 interfaces on each ASA, for a total of 32 interfaces in the spanned EtherChannel.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

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

Step 4 (Optional) To override the media type, duplex, speed, and pause frames for flow control for all member interfaces, click **Configure Hardware Properties**. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

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

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

- In the **MAC Address Cloning** area, set a manual MAC address for the EtherChannel. Do not set the Standby MAC Address; it is ignored. You must configure a MAC address for a Spanned EtherChannel so that the MAC address does not change when the current master unit leaves the cluster; with a manually-configured MAC address, the MAC address stays with the current master unit.

In multiple context mode, if you share an interface between contexts, auto-generation of MAC addresses is enabled by default, so that you only need to set the MAC address manually for a shared interface if you disable auto-generation. Note that you must manually configure the MAC address for non-shared interfaces. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

- (Optional) If you are connecting the ASA to two switches in a VSS or vPC, then you should enable VSS load balancing by checking the **Enable load balancing between switch pairs in VSS or vPC** mode check box. This feature ensures that the physical link connections between the ASAs to the VSS (or vPC) pair are balanced.

In the **Member Interface Configuration** area, you must then identify to which switch a given interface is connected, **1** or **2**.



Note We recommend that you do *not* set the **Minimum Active Members** and the **Maximum Active Members**.

Step 6 (Optional) Configure VLAN subinterfaces on this EtherChannel. The rest of this procedure applies to the subinterfaces.

Step 7 (Multiple context mode) Before you complete this procedure, you need to allocate interfaces to contexts.

- a. Click **OK** to accept your changes.

- b. Allocate interfaces.
- c. Change to the context that you want to configure: in the **Device List** pane, double-click the context name under the active device IP address.
- d. Choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane, select the port-channel interface that you want to customize, and click **Edit**.

The **Edit Interface** dialog box appears.

- Step 8** Click the **General** tab.
- Step 9** (Transparent Mode) From the **Bridge Group** drop-down list, choose the bridge group to which you want to assign this interface.
- Step 10** In the **Interface Name** field, enter a name up to 48 characters in length.
- Step 11** In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).
- Step 12** (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.
- Step 13** (Routed Mode) To configure an IPv6 address, click the **IPv6** tab.
- For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.
- a. Check the **Enable IPv6** check box.
 - b. In the **Interface IPv6 Addresses** area, click **Add**.
- The **Add IPv6 Address for Interface** dialog box appears.
- Note:** The **Enable address autoconfiguration** option is not supported.
- c. In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:DB8::BA98:0:3210/64.
 - d. (Optional) To use the Modified EUI-64 interface ID as the host address, check the **EUI-64** check box. In this case, just enter the prefix in the **Address/Prefix Length** field.
 - e. Click **OK**.
- Step 14** Click **OK** to return to the **Interfaces** screen.
- Step 15** Click **Apply**.
-

Related Topics

- [Configure the Cluster Interface Mode on the Master Unit, page 8-38](#)
- [Configure Bridge Groups, page 13-9](#)
- [Create or Join an ASA Cluster, page 8-47](#)
- [Configure an EtherChannel, page 12-8](#)
- [Guidelines for EtherChannels and Redundant Interfaces, page 12-4](#)
- [Connecting to a VSS or vPC, page 8-14](#)
- [Enable the Physical Interface and Configure Ethernet Parameters, page 9-6](#)
- [Configure VLAN Subinterfaces and 802.1Q Trunking, page 13-3](#)
- [Configure a Security Context, page 8-18](#)

- [Security Levels, page 13-2](#)
- [Configure ASA Cluster Parameters, page 8-50](#)
- [Guidelines for ASA Clustering, page 8-32](#)

Create or Join an ASA Cluster

Each unit in the cluster requires a bootstrap configuration to join the cluster. Run the High Availability and Scalability wizard on one unit (that will become the master unit) to create the cluster, and then add slave units to it.



Note

For the master unit, if you want to change the default of the cLACP system ID and priority, you cannot use the wizard; you must configure the cluster manually.

Before You Begin

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

Procedure

-
- Step 1** Choose **Wizards > High Availability and Scalability Wizard**. See select wizard guidelines in the following steps.
- Step 2** On the **Interfaces** screen, you cannot create new EtherChannels from this screen (except for the cluster control link).
- Step 3** On the ASA Cluster Configuration screen, configure bootstrap settings including:
- **Member Priority**—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.
 - (Optional) **Shared Key**—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
 - (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded

devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.



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

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster health check feature, which includes unit health monitoring and interface health monitoring. To determine unit health, the ASA cluster units send keepalive messages on the cluster control link to other units. If a unit does not receive any keepalive messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead. The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular unit, but there are active ports under the same logical interface on other units, then the unit is removed from the cluster. If a unit does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. Health check is enabled by default for all interfaces. You can optionally later disable it per interface.



Note When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you must disable the health check. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check.

- **Time to Wait Before Device Considered Failed**—This value determines the amount of time between unit keepalive status messages, between .8 and 45 seconds; The default is 3 seconds. Note that the holdtime value only affects the *unit* health check; for interface health, the ASA uses the interface status (up or down).
- (Optional) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.
- (Optional) **Replicate console output to the master's console**—Enables console replication from slave units to the master unit. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, slave units send the console messages to the master unit so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.
- **Cluster Control Link**—Specifies the cluster control link interface.
 - (Optional) **MTU**—Specifies the maximum transmission unit for the cluster control link interface, between 64 and 65,535 bytes. Data that is larger than the MTU value is fragmented before being sent. The default MTU is 1500 bytes. If you already enabled jumbo frame

reservation, we suggest setting the MTU to 1600 bytes or greater. If you want to use jumbo frames and have not pre-enabled jumbo frame reservation, you should quit the wizard, enable jumbo frames, and then restart this procedure.

Step 4 Click **Finish**.

Step 5 The ASA scans the running configuration for incompatible commands for features that are not supported with clustering, including commands that may be present in the default configuration. Click **OK** to delete the incompatible commands. If you click **Cancel**, then clustering is not enabled.

Step 6 After a period of time while ASDM enables clustering and reconnects to the ASA, the Information screen appears confirming that the ASA was added to the cluster.




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

Step 7 To add a slave unit, click **Yes**.

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

Step 8 In the **Deployment Options** area, choose one of the following **Deploy By** options:

- **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the slave (temporary) management IP address. Enter the slave management IP address, username, and password.
- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the slave unit CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Related Topics

- [Configure ASA Cluster Parameters, page 8-50](#)
- [Enable Jumbo Frame Support, page 11-8](#)
- [Configure Spanned EtherChannels, page 8-44](#)
- [Configure Individual Interfaces \(Recommended for the Management Interface\), page 8-41](#)

- [Interface Monitoring, page 8-9](#)

Manage ASA Cluster Members

After you deploy the cluster, you can change the configuration and manage cluster members.

- [Configure ASA Cluster Parameters, page 8-50](#)
- [Add a New Slave from the Master Unit, page 8-53](#)
- [Become an Inactive Member, page 8-54](#)
- [Inactivate a Slave Member from the Master Unit, page 8-55](#)
- [Leave the Cluster, page 8-56](#)
- [Change the Master Unit, page 8-57](#)
- [Execute a Command Cluster-Wide, page 8-58](#)

Configure ASA Cluster Parameters

If you do not use the wizard to add a unit to the cluster, you can configure the cluster parameters manually. If you already enabled clustering, you can edit some cluster parameters; others that cannot be edited while clustering is enabled are grayed out. This procedure also includes advanced parameters that are not included in the wizard.

Before You Begin

- Pre-configure the cluster control link interfaces on each unit before joining the cluster. For a single interface, you must enable it; do not configure any other settings. For an EtherChannel interface, enable it and set the EtherChannel mode to On.
- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

Procedure

-
- Step 1** Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

If your device is already in the cluster, and is the master unit, then this pane is on the Cluster Configuration tab.

- Step 2** Check the **Configure ASA cluster settings** check box.

If you uncheck the check box, the settings are erased. Do not check **Participate in ASA cluster** until after you have set all your parameters.



Note

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

- Step 3** Configure the following bootstrap parameters:

- **Cluster Name**—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per unit. All members of the cluster must use the same name.
- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.
- (Optional) **Shared Key**—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.
- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster health check feature, which includes unit health monitoring and interface health monitoring. **Note:** When you are adding new units to the cluster, and making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete. You can re-enable this feature after cluster and topology changes are complete. To determine unit health, the ASA cluster units send keepalive messages on the cluster control link to other units. If a unit does not receive any keepalive messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead. Interface status messages detect link failure. If all physical ports for a given logical interface fail on a particular unit, but there are active ports under the same logical interface on other units, then the unit is removed from the cluster. If a unit does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. Health check is enabled by default for all interfaces. You can optionally disable it per interface on the **Cluster Interface Health Monitoring** tab.

**Note**

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you must disable the health check. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check.

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

- (Optional) **Replicate console output to the master's console**—Enables console replication from slave units to the master unit. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, slave units send the console messages to the master unit so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.
- **Cluster Control Link**—Specifies the cluster control link interface. This interface cannot have a name configured; available interfaces are shown in the drop-down list.
 - **Interface**—Specifies the interface ID, preferably an EtherChannel. Subinterfaces and Management type interfaces are not allowed.
 - **IP Address**—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.
 - **Subnet Mask**—Specifies the subnet mask.
 - (Optional) **MTU**—Specifies the maximum transmission unit for the cluster control link interface, between 64 and 65,535 bytes. Data that is larger than the MTU value is fragmented before being sent. The default MTU is 1500 bytes. We suggest setting the MTU to 1600 bytes or greater, which requires you to enable jumbo frame reservation.
- (Optional) **Cluster LACP**—When using Spanned EtherChannels, the ASA uses cLACP to negotiate the EtherChannel with the neighbor switch. ASAs in a cluster collaborate in cLACP negotiation so that they appear as a single (virtual) device to the switch.
 - **Enable static port priority**—Disables dynamic port priority in LACP. Some switches do not support dynamic port priority, so this parameter improves switch compatibility. Moreover, it enables support of more than 8 active spanned EtherChannel members, up to 32 members. Without this parameter, only 8 active members and 8 standby members are supported. If you enable this parameter, then you cannot use any standby members; all members are active. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units.
 - **Virtual System MAC Address**—Sets the cLACP system ID, which is in the format of a MAC address. All ASAs use the same system ID: auto-generated by the master unit (the default) and replicated to all slaves; or manually specified in the form *H.H.H*, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units. However, you cannot change this value after you enable clustering.
 - **System Priority**—Sets the system priority, between 1 and 65535. The priority is used to decide which unit is in charge of making a bundling decision. By default, the ASA uses priority 1, which is the highest priority. The priority needs to be higher than the priority on the switch. This parameter is not part of the bootstrap configuration, and is replicated from the master unit to the slave units. However, you cannot change this value after you enable clustering.

Step 4 Check the **Participate in ASA cluster** check box to join the cluster.

Step 5 (Optional) Click the **Cluster Interface Health Monitoring** tab to disable health monitoring on one or more interfaces. Select an interface in the **Monitored Interfaces** box, and click **Add** to move it to the **Unmonitored Interfaces** box.

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can monitor any port-channel ID, redundant ID, or single physical interface ID. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Step 6 Click **Apply**.

Related Topics

- [Interface Monitoring, page 8-9](#)
- [Enable Jumbo Frame Support, page 11-8](#)

Add a New Slave from the Master Unit

You can add additional slaves to the cluster from the master unit. You can also add slaves using the High Availability and Scalability wizard. Adding a slave from the master unit has the benefit of configuring the cluster control link and setting the cluster interface mode on each slave unit you add.


You can alternatively log into the slave unit and configure clustering directly on the unit. However, after you enable clustering, your ASDM session will be disconnected, and you will have to reconnect.

Before You Begin

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.
- If you want to send the bootstrap configuration over the management network, be sure the slave unit has an accessible IP address.

Procedure

- Step 1** Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Members**.
- Step 2** Click **Add**.
- Step 3** Configure the following parameters:
- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
 - **Member Priority**—Sets the priority of this unit for master unit elections, between 1 and 100, where 1 is the highest priority.
 - **Cluster Control Link > IP Address**—Specifies a unique IP address for this member for the cluster control link, on the same network as the master cluster control link.
 - In the **Deployment Options** area, choose one of the following **Deploy By** options:
 - **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the slave (temporary) management IP address. Enter the slave management IP address, username, and password.
 - **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the slave unit CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Step 4 Click **OK**, then **Apply**.

Related Topics

- [Configure ASA Cluster Parameters, page 8-50](#)

Become 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. 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.

Before You Begin

- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

- Step 1** Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.
If your device is already in the cluster, and is the master unit, then this pane is on the Cluster Configuration tab.
- Step 2** Uncheck the **Participate in ASA cluster** check box.

**Note**

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

Step 3 Click **Apply**.

If this unit was the master unit, a new master election takes place, and a different member becomes the master unit.

The cluster configuration is maintained, so that you can enable clustering again later.

Related Topics

- [Leave the Cluster, page 8-56](#)

Inactivate a Slave Member from the Master Unit

To inactivate a slave member, perform the following steps.

**Note**

When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the unit altogether from the cluster. 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.

Before You Begin

For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure**Step 1** Remove the unit from the cluster:

```
cluster remove unit unit_name
```

Example:

```
ciscoasa(config)# cluster remove unit ?
```

```
Current active units in the cluster:  
asa2
```

```
ciscoasa(config)# cluster remove unit asa2
```

```
WARNING: Clustering will be disabled on unit asa2. To bring it back  
to the cluster please logon to that unit and re-enable clustering
```

The bootstrap configuration remains intact, as well as the last configuration synced from the master unit, so that 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.

To view member names, enter **cluster remove unit ?**, or enter the **show cluster info** command.

Step 1 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

Step 2 Select the slave that you want to remove, and click **Delete**.

The slave bootstrap configuration remains intact, so that you can later re-add the slave without losing your configuration.

Step 3 Click **Apply**.

Related Topics

- [Leave the Cluster, page 8-56](#)

Leave 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.

Before You Begin

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.

Procedure

Step 1 For a slave unit, disable clustering:

```
cluster group cluster_name
no enable
```

Example:

```
ciscoasa(config)# cluster group cluster1
ciscoasa(cfg-cluster)# no enable
```

You cannot make configuration changes while clustering is enabled on a slave unit.

Step 2 Clear the cluster configuration:

```
clear configure cluster
```

The ASA shuts down all interfaces including the management interface and cluster control link.

Step 3 Disable cluster interface mode:

```
no 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 the backup configuration to the running configuration:

```
copy backup_cfg running-config
```

Example:

```
ciscoasa(config)# copy backup_cluster.cfg running-config
```

```
Source filename [backup_cluster.cfg]?
```

```
Destination filename [running-config]?
```

```
ciscoasa(config)#
```

Step 5 Save the configuration to startup:

```
write memory
```

Step 6 If you do not have a backup configuration, reconfigure management access. Be sure to change the interface IP addresses, and restore the correct hostname, for example.

Related Topics

- [Chapter 2, “Getting Started.”](#)

Change the Master Unit



Caution

The best method to change the master unit is to disable clustering on the master unit, 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.

To change the master unit, perform the following steps.

Before You Begin

For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

- Step 1** Choose **Monitoring > ASA Cluster > Cluster Summary**.
- Step 2** From the **Change Master To** drop-down list, choose a slave unit to become master, and click **Make Master**.
- Step 3** You are prompted to confirm the master unit change. Click **Yes**.
- Step 4** Quit ASDM, and reconnect using the Main cluster IP address.

Related Topics

- [Become an Inactive Member, page 8-54](#)
- [Centralized Features for Clustering, page 8-24](#)

Execute 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.

Before You Begin

Perform this procedure at the Command Line Interface tool: choose **Tools > Command Line Interface**.

Procedure

- Step 1** Send a command to all members, or if you specify the unit name, a specific member:

```
cluster exec [unit unit_name] command
```

Example:

```
cluster exec show xlate
```

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:

```
cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each unit, are copied to the TFTP server. The destination capture file name is automatically attached with the unit name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster unit names.

The following sample output for the **cluster exec show port-channel** summary command shows EtherChannel information for each member in the cluster:

```
cluster exec show port-channel summary
primary(LOCAL):*****
Number of channel-groups in use: 2
Group Port-channel Protocol Span-cluster Ports
-----+-----+-----+-----+-----
1      Po1          LACP      Yes   Gi0/0(P)
2      Po2          LACP      Yes   Gi0/1(P)
secondary:*****
Number of channel-groups in use: 2
Group Port-channel Protocol Span-cluster Ports
-----+-----+-----+-----+-----
1      Po1          LACP      Yes   Gi0/0(P)
2      Po2          LACP      Yes   Gi0/1(P)
```

Monitoring the ASA Cluster

You can monitor and troubleshoot cluster status and connections.

- [Monitoring Cluster Status, page 8-59](#)

- [Capturing Packets Cluster-Wide](#), page 8-59
- [Monitoring Cluster Resources](#), page 8-59
- [Monitoring Cluster Traffic](#), page 8-60
- [Monitoring the Cluster Control Link](#), page 8-60
- [Configuring Logging for Clustering](#), page 8-60

Monitoring Cluster Status

See the followingscreens for monitoring cluster status:

- **Monitoring > ASA Cluster > Cluster Summary**

This pane shows cluster information about the unit to which you are connected, as well as other units in the cluster. You can also change the master unit from this pane.

- **Cluster Dashboard**

On the home page on the master unit, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard.

Related Topics

- [Cluster Dashboard Tab](#), page 3-23
- [Cluster Firewall Dashboard Tab](#), page 3-24

Capturing Packets Cluster-Wide

See the followingscreens for capturing packets in a cluster:

Wizards > Packet Capture Wizard

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the master unit, which is then automatically enabled on all of the slave units in the cluster.

Related Topics

- [Configure and Run Captures with the Packet Capture Wizard](#), page 37-1

Monitoring Cluster Resources

See the followingscreens for monitoring cluster resources:

- **Monitoring > ASA Cluster > System Resources Graphs > CPU**

This pane lets you create graphs or tables showing the CPU utilization across the cluster members.

- **Monitoring > ASA Cluster > System Resources Graphs > Memory**. This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster members.

Monitoring Cluster Traffic

See the followingscreens for monitoring cluster traffic:

- **Monitoring > ASA Cluster > Traffic Graphs > Connections.**

This pane lets you create graphs or tables showing the Connections across the cluster members.

- **Monitoring > ASA Cluster > Traffic Graphs > Throughput.**

This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

Monitoring the Cluster Control Link

See the following screens for monitoring cluster status:

Monitoring > Properties > System Resources Graphs > Cluster Control Link.

This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Configuring Logging for Clustering

See the followingscreens for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

Each unit in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster.

Related Topics

- [Include the Device ID in Non-EMBLEM Format Syslog Messages, page 38-21](#)

Examples for ASA Clustering

These examples include all cluster-related ASA configuration for typical deployments.

- [Sample ASA and Switch Configuration, page 8-60](#)
- [Firewall on a Stick, page 8-63](#)
- [Traffic Segregation, page 8-65](#)
- [Spanned EtherChannel with Backup Links \(Traditional 8 Active/8 Standby\), page 8-67](#)

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 8-61](#)
- [Cisco IOS Switch Configuration, page 8-62](#)

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

```

Cisco IOS Switch Configuration

```

interface GigabitEthernet1/0/15
  switchport access vlan 201
  switchport mode access
  spanning-tree portfast
  channel-group 10 mode active
!
interface GigabitEthernet1/0/16
  switchport access vlan 201
  switchport mode access
  spanning-tree portfast
  channel-group 10 mode active
!
interface GigabitEthernet1/0/17
  switchport access vlan 401
  switchport mode access
  spanning-tree portfast
  channel-group 11 mode active
!
interface GigabitEthernet1/0/18

```



```


switchport access vlan 401
switchport mode access
spanning-tree portfast
channel-group 11 mode active

interface Port-channel10
switchport access vlan 201
switchport mode access

interface Port-channel11
switchport access vlan 401
switchport mode access

```

Firewall on a Stick



Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each ASA has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The ASA is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If an ASA becomes unavailable, the switch will rebalance traffic between the remaining units.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/8
  no shutdown
  description CCL

cluster group cluster1
  local-unit asa1
  cluster-interface tengigabitethernet0/8 ip 192.168.1.1 255.255.255.0
  priority 1
  key chuntheunavoidable
  enable noconfirm
```

ASA2 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/8
  no shutdown
  description CCL

cluster group cluster1
  local-unit asa2
  cluster-interface tengigabitethernet0/8 ip 192.168.1.2 255.255.255.0
  priority 2
  key chuntheunavoidable
  enable as-slave
```

ASA3 Slave Bootstrap Configuration

```
interface tengigabitethernet 0/8
  no shutdown
  description CCL

cluster group cluster1
  local-unit asa3
  cluster-interface tengigabitethernet0/8 ip 192.168.1.3 255.255.255.0
  priority 3
  key chuntheunavoidable
  enable as-slave
```

Master Interface Configuration

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

interface management 0/0
  nameif management
  ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
  ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
  security-level 100
  management-only
  no shutdown


interface tengigabitethernet 0/9
  channel-group 2 mode active
  no shutdown
interface port-channel 2
  port-channel span-cluster
interface port-channel 2.10
```

```

vlan 10
nameif inside
ip address 10.10.10.5 255.255.255.0
ipv6 address 2001:DB8:1::5/64
mac-address 000C.F142.4CDE
interface port-channel 2.20
vlan 20
nameif outside
ip address 209.165.201.1 255.255.255.224
ipv6 address 2001:DB8:2::8/64
mac-address 000C.F142.5CDE

```

Traffic Segregation



You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Master Bootstrap Configuration

```
interface tengigabitethernet 0/6
channel-group 1 mode on
```

```

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

cluster group cluster1
  local-unit asa1
  cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
  priority 1
  key chuntheunavoidable
  enable noconfirm

```

ASA2 Slave Bootstrap Configuration

```

interface tengigabitethernet 0/6
  channel-group 1 mode on
  no shutdown
interface tengigabitethernet 0/7
  channel-group 1 mode on
  no shutdown
interface port-channel 1
  description CCL

cluster group cluster1
  local-unit asa2
  cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
  priority 2
  key chuntheunavoidable
  enable as-slave

```

ASA3 Slave Bootstrap Configuration

```

interface tengigabitethernet 0/6
  channel-group 1 mode on
  no shutdown
interface tengigabitethernet 0/7
  channel-group 1 mode on
  no shutdown
interface port-channel 1
  description CCL

cluster group cluster1
  local-unit asa3
  cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
  priority 3
  key chuntheunavoidable
  enable as-slave

```

Master Interface Configuration

```

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

interface management 0/0
  nameif management
  ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
  ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
  security-level 100
  management-only

```


```
no shutdown

interface tengigabitethernet 0/8
  channel-group 2 mode active
  no shutdown
interface port-channel 2
  port-channel span-cluster
  nameif inside
  ip address 10.10.10.5 255.255.255.0
  ipv6 address 2001:DB8:1::5/64
  mac-address 000C.F142.4CDE

interface tengigabitethernet 0/9
  channel-group 3 mode active
  no shutdown
interface port-channel 3
  port-channel span-cluster
  nameif outside
  ip address 209.165.201.1 255.255.255.224
  ipv6 address 2001:DB8:2::8/64
  mac-address 000C.F142.5CDE
```


Spanned EtherChannel with Backup Links (Traditional 8 Active/8 Standby)

The maximum number of active ports in a traditional 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

```

History for ASA Clustering

Feature Name	Platform Releases	Feature Information
ASA Clustering for the ASA 5580 and 5585-X	9.0(1)	ASA Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. ASA clustering is supported for the ASA 5580 and the ASA 5585-X; all units in a cluster must be the same model with the same hardware specifications. See the configuration guide for a list of unsupported features when clustering is enabled. We introduced or modified the following screens: - Home > Device Dashboard - Home > Cluster Dashboard - Home > Cluster Firewall Dashboard - Configuration > Device Management > Advanced > Address Pools > MAC Address Pools - Configuration > Device Management > High Availability and Scalability > ASA Cluster - Configuration > Device Management > Logging > Syslog Setup > Advanced - Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > Advanced - Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > IPv6 - Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface > Advanced - Configuration > Firewall > Advanced > Per-Session NAT Rules - Monitoring > ASA Cluster - Monitoring > Properties > System Resources Graphs > Cluster Control Link - Tools > Preferences > General - Tools > System Reload - Tools > Upgrade Software from Local Computer - Wizards > High Availability and Scalability Wizard - Wizards > Packet Capture Wizard - Wizards > Startup Wizard
ASA 5500-X support for clustering	9.1(4)	The ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X now support 2-unit clusters. Clustering for 2 units is enabled by default in the base license; for the ASA 5512-X, you need the Security Plus license. We did not modify any ASDM screens.

Feature Name	Platform Releases	Feature Information
Improved VSS and vPC support for health check monitoring	9.1(4)	If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, you can now increase stability with health check monitoring. For some switches, such as the Cisco Nexus 5000, when one unit in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends keepalive messages on one of these EtherChannel interfaces. When you enable the VSS/vPC health check feature, the ASA floods the keepalive messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster
Support for cluster members at different geographical locations (inter-site); Individual Interface mode only	9.1(4)	You can now place cluster members at different geographical locations when using Individual Interface mode. We did not modify any ASDM screens.
Support for cluster members at different geographical locations (inter-site) for transparent mode	9.2(1)	You can now place cluster members at different geographical locations when using Spanned EtherChannel mode in transparent firewall mode. Inter-site clustering with spanned EtherChannels in routed firewall mode is not supported. We did not modify any ASDM screens.
Static LACP port priority support for clustering	9.2(1)	Some switches do not support dynamic port priority with LACP (active and standby links). You can now disable dynamic port priority to provide better compatibility with spanned EtherChannels. You should also follow these guidelines: • Network elements on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped. • Port-channel bundling downtime should not exceed the configured keepalive interval. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster

Feature Name	Platform Releases	Feature Information
Support for 32 active links in a spanned EtherChannel	9.2(1)	ASA EtherChannels now support up to 16 active links. With <i>spanned</i> EtherChannels, that functionality is extended to support up to 32 active links across the cluster when used with two switches in a vPC and when you disable dynamic port priority. The switches must support EtherChannels with 16 active links, for example, the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module. For switches in a VSS or vPC that support 8 active links, you can now configure 16 active links in the spanned EtherChannel (8 connected to each switch). Previously, the spanned EtherChannel only supported 8 active links and 8 standby links, even for use with a VSS/vPC. Note If you want to use more than 8 active links in a spanned EtherChannel, you cannot also have standby links; the support for 9 to 32 active links requires you to disable cLACP dynamic port priority that allows the use of standby links. We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster
Support for 16 cluster members for the ASA 5585-X	9.2(1)	The ASA 5585-X now supports 16-unit clusters. We did not modify any screens.
BGP support for ASA clustering	9.3(1)	We added support for BGP with ASA clustering. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > General
Inter-site deployment in transparent mode with the ASA cluster firewalling between inside networks	9.3(2)	You can now deploy a cluster in transparent mode between inside networks and the gateway router at each site (AKA East-West insertion), and extend the inside VLANs between sites. We recommend using Overlay Transport Virtualization (OTV), but you can use any method that ensures that the overlapping MAC Addresses and IP addresses of the gateway router do not leak between sites. Use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide the same virtual MAC and IP addresses to the gateway routers.
Enable and disable ASA cluster health monitoring per interface	9.4(1)	You can now enable or disable health monitoring per interface. Health monitoring is enabled by default on all port-channel, redundant, and single physical interfaces. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored. You might want to disable health monitoring of non-essential interfaces, for example, the management interface. We introduced the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Interface Health Monitoring
ASA clustering support for DHCP relay	9.4(1)	You can now configure DHCP relay on the ASA cluster. Client DHCP requests are load-balanced to the cluster members using a hash of the client MAC address. DHCP client and server functions are still not supported. We did not modify any screens.
SIP inspection support in ASA clustering	9.4(1)	You can now configure SIP inspection on the ASA cluster. A control flow can be created on any unit (due to load balancing), but its child data flows must reside on the same unit. TLS Proxy configuration is not supported. We did not modify any screens.



PART 3

Interfaces



Basic Interface Configuration

This chapter includes basic interface configuration including Ethernet settings and Jumbo frame configuration.



Note

For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.



Note

For the ASA Services Module interfaces, see the ASA Services Module quick start guide.

- [About Basic Interface Configuration, page 11-1](#)
- [Licensing for Basic Interface Configuration, page 11-4](#)
- [Guidelines for Basic Interface Configuration, page 11-5](#)
- [Default Settings for Basic Interface Configuration, page 11-6](#)
- [Enable the Physical Interface and Configure Ethernet Parameters, page 11-6](#)
- [Examples for Basic Interfaces, page 11-9](#)
- [History for Basic Interface Configuration, page 11-9](#)

About Basic Interface Configuration

This section describes interface features and special interfaces.

- [Auto-MDI/MDIX Feature, page 11-1](#)
- [Management Interface, page 11-2](#)

Auto-MDI/MDIX Feature

For RJ-45 interfaces, 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.

Management Interface

The management interface, depending on your model, is a separate interface just for management traffic.

- [Management Interface Overview, page 11-2](#)
- [Management Slot/Port Interface, page 11-2](#)
- [Use Any Interface for Management-Only Traffic, page 11-3](#)
- [Management Interface for Transparent Mode, page 11-3](#)
- [No Support for Redundant Management Interfaces, page 11-4](#)
- [Management Interface on All Models Except for the ASA 5585-X, page 11-4](#)

Management Interface Overview

You can manage the ASA by connecting to:

- Any through-traffic interface
- A dedicated Management *Slot/Port* interface (if available for your model)

You may need to configure management access to the interface according to [Chapter 34, “Management Access.”](#)

Management *Slot/Port* Interface

The following table shows the Management interfaces per model.

Table 11-1 Management Interfaces Per Model

Model	Management 0/0	Management 0/1	Management 1/0	Management 1/1	Configurable for Through Traffic	Subinterfaces Allowed
ASA 5506-X	No	No	No	Yes	No	No
ASA 5508-X	No	No	No	Yes	No	No
ASA 5512-X	Yes	No	No	No	No	No
ASA 5515-X	Yes	No	No	No	No	No
ASA 5516-X	No	No	No	Yes	No	No
ASA 5525-X	Yes	No	No	No	No	No
ASA 5545-X	Yes	No	No	No	No	No
ASA 5555-X	Yes	No	No	No	No	No

Table 11-1 **Management Interfaces Per Model**

Model	Management 0/0	Management 0/1	Management 1/0	Management 1/1	Configurable for Through Traffic	Subinterfaces Allowed
ASA 5585-X	Yes	Yes	Yes 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.	Yes	Yes	Yes
ASASM	No	No	No	No	N/A	N/A
ASAv	Yes	No	No	No	No	No

**Note**

If you installed a module, then the module management interface(s) provides management access for the module only. For models with software modules, the software module uses the same physical Management interface as the ASA.

Use Any Interface for Management-Only Traffic

You can use any interface as a dedicated management-only interface by configuring it for management traffic, including an EtherChannel interface.

Management Interface for Transparent Mode

In transparent firewall mode, in addition to the maximum allowed through-traffic interfaces, you can also use the Management interface (either the physical interface, a subinterface (if supported for your model), or an EtherChannel interface comprised of Management interfaces (if you have multiple Management interfaces)) as a separate management interface. You cannot use any other interface types as management interfaces.

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 5555-X and lower do not allow subinterfaces on the Management interface, so for per-context management, you must connect to a data interface.

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 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 Interface on All Models Except for the ASA 5585-X

The Management interface has the following characteristics:

- No through traffic support
- No subinterface support
- No priority queue support
- No multicast MAC support
- The software module shares the Management interface. Separate MAC addresses and IP addresses are supported for the ASA and module. You must perform configuration of the module IP address within the module operating system. However, physical characteristics (such as enabling the interface) are configured on the ASA.

Licensing for Basic Interface Configuration

Model	License Requirement
ASA 5506-X series	Interfaces of all types: Base License: 536 Security Plus License: 636
ASA 5508-X	Interfaces of all types: Base License: 716
ASA 5512-X	Interfaces of all types: Base License: 716 Security Plus License: 916
ASA 5515-X	Interfaces of all types: Base License: 916
ASA 5516-X	Interfaces of all types: Base License: 716
ASA 5525-X	Interfaces of all types: Base License: 1316
ASA 5545-X	Interfaces of all types: Base License: 1716

Model	License Requirement
ASA 5555-X	Interfaces of all types: Base License: 2516
ASA 5585-X	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: Base and Security Plus License: 4612
ASAv5 and ASAv10	Interfaces of all types: Standard and Premium License: 716
ASAv30	Interfaces of all types: Standard and Premium License: 1316

**Note**

Interfaces of all types comprise the maximum number of combined interfaces; for example, VLANs, VXLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** defined in the configuration counts against this limit.

Guidelines for Basic Interface Configuration

Firewall Mode

For multiple context, transparent mode, each context must use different interfaces; you cannot share an interface across contexts.

Failover

You cannot share a failover or state interface with a data interface.

Additional Guidelines

Some management-related services are not available until a non-management interface is enabled, and the ASA achieves a “System Ready” state. The ASA generates the following syslog message when it is in a “System Ready” state:

```
%ASA-6-199002: Startup completed. Beginning operation.
```

Default Settings for Basic Interface Configuration

This section lists default settings for interfaces if you do not have a factory default configuration.

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.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- VXLAN VNI interfaces—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.
- For fiber interfaces for the 5585-X, the speed is set for automatic link negotiation.

Default Connector Type

Some models 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.

Enable the Physical Interface and Configure 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

Before You Begin

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

- Step 1** Depending on your context mode:
- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

By default, all physical interfaces are listed.

- Step 2** Click a physical interface that you want to configure, and click **Edit**.
The **Edit Interface** dialog box appears.



Note In single mode, this procedure only covers a subset of the parameters on the **Edit Interface** dialog box. Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts.

- Step 3** To enable the interface, check the **Enable Interface** check box.

- Step 4** To add a description, enter text in the Description field.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

- Step 5** (Optional) To set the media type, duplex, speed, and enable pause frames for flow control, click **Configure Hardware Properties**.

- Depending on the interface type, you can choose either **RJ-45** or **SFP** from the **Media Type** drop-down list.
RJ-45 is the default.
- To set the duplex for RJ-45 interfaces, choose **Full**, **Half**, or **Auto**, depending on the interface type, from the **Duplex** drop-down list.



Note The duplex setting for an EtherChannel interface must be **Full** or **Auto**.

- To set the speed, choose a value from the **Speed** drop-down list.

The speeds available depend on the interface type. For SFP interfaces, you can set the speed to **Negotiate** or **Nonegotiate**. **Negotiate** (the default) enables link negotiation, which exchanges flow-control parameters and remote fault information. **Nonegotiate** does not negotiate link parameters. For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature.

- To enable pause (XOFF) frames for flow control on Gigabit and Ten Gigabit Ethernet interfaces, check the **Enable Pause Frame** check box.

If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. Pause (XOFF) and XON frames are generated automatically by the NIC hardware based on the FIFO buffer usage. A pause frame is sent when the buffer usage exceeds the high-water mark. The default *high_water* value is 128 KB (10 GigabitEthernet) and 24 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and 47 KB (1 GigabitEthernet). After a pause is sent, an XON frame can be sent when the buffer usage is

reduced below the low-water mark. By default, the *low_water* value is 64 KB (10 GigabitEthernet) and 16 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and 47 KB (1 GigabitEthernet). The link partner can resume traffic after receiving an XON, or after the XOFF expires, as controlled by the timer value in the pause frame. The default *pause_time* value is 26624; you can set it between 0 and 65535. If the buffer usage is consistently above the high-water mark, pause frames are sent repeatedly, controlled by the pause refresh threshold value.

To change the default values for the Low Watermark, High Watermark, and Pause Time, uncheck the **Use Default Values** check box.



Note Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

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

Step 6 Click **OK** to accept the **Interface** changes.

Enable Jumbo Frame Support

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.

Before You Begin

- 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 9198. In multiple context mode, set the MTU within each context.
- Be sure to adjust the TCP MSS, either to disable it for non-VPN traffic, or to increase it in accord with the MTU.

Procedure

Step 1 Depending on your context mode:

- Multiple mode—To enable jumbo frame support, choose **Configuration > Context Management > Interfaces**, and click the **Enable jumbo frame support** check box.
- Single mode—Setting the MTU larger than 1500 bytes automatically enables jumbo frames. To manually enable or disable this setting, choose **Configuration > Device Setup > Interface Settings > Interfaces**, and click the **Enable jumbo frame support** check box.

Related Topics

- [About Advanced Interface Configuration, page 16-1](#)

- [Configure the MAC Address, MTU, and TCP MSS, page 16-5](#)

Examples for Basic Interfaces

See the following configuration examples.

- [Physical Interface Parameters Example, page 11-9](#)
- [Multiple Context Mode Example, page 11-9](#)

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
```

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
```

History for Basic Interface Configuration

Table 11-2 History for Interfaces

Feature Name	Releases	Feature Information
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.
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.

Table 11-2 *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. This feature is also supported on the ASA 5585-X. We modified the following screen: Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > Advanced.
Support for Pause Frames for Flow Control on the ASA 5580 Ten Gigabit Ethernet Interfaces	8.2(2)	You can now enable pause (XOFF) frames for flow control. This feature is also supported on the ASA 5585-X. We modified the following screens: (Single Mode) Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Support for Pause Frames for Flow Control on Gigabit Ethernet Interfaces	8.2(5)/8.4(2)	You can now enable pause (XOFF) frames for flow control for Gigabit Ethernet interfaces on all models. We modified the following screens: (Single Mode) Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.



EtherChannel and Redundant Interfaces

This chapter tells how to configure EtherChannels and redundant interfaces.



Note

For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.



Note

For ASA cluster interfaces, which have special requirements, see [Chapter 8, “ASA Cluster.”](#)

- [About EtherChannels and Redundant Interfaces, page 12-1](#)
- [Guidelines for EtherChannels and Redundant Interfaces, page 12-4](#)
- [Default Settings for EtherChannels and Redundant Interfaces, page 12-6](#)
- [Configure a Redundant Interface, page 12-6](#)
- [Configure an EtherChannel, page 12-8](#)
- [Examples for EtherChannel and Redundant Interfaces, page 12-12](#)
- [History for EtherChannels and Redundant Interfaces, page 12-12](#)

About EtherChannels and Redundant Interfaces

This section describes EtherChannels and Redundant Interfaces.

- [Redundant Interfaces, page 12-1](#)
- [EtherChannels, page 12-2](#)

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. When the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

Related Topics

- [Configure the MAC Address, MTU, and TCP MSS, page 16-5](#)
- [Configure Multiple Contexts, page 8-15](#)

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.

- [Channel Group Interfaces, page 12-2](#)
- [Connecting to an EtherChannel on Another Device, page 12-2](#)
- [Link Aggregation Control Protocol, page 12-3](#)
- [Load Balancing, page 12-4](#)
- [EtherChannel MAC Address, page 12-4](#)

Channel Group Interfaces

Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure. For 16 active interfaces, be sure that your switch supports the feature (for example, the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).


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 interface 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 or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch.

Figure 12-1 Connecting to a VSS/vPC

If you use the ASA in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each ASA. 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 12-2 Active/Standby Failover and VSS/vPC

Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDU)s 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). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value mod active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

For a spanned EtherChannel in clustering, load balancing occurs on a per ASA basis. For example, if you have 32 active interfaces in the spanned EtherChannel across 8 ASAs, with 4 interfaces per ASA in the EtherChannel, then load balancing only occurs across the 4 interfaces on the ASA.

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.

Related Topics

- [Customize the EtherChannel, page 12-10](#)

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.

Guidelines for EtherChannels and Redundant Interfaces

Failover

- When you use a redundant or EtherChannel interface as a failover link, it must be pre-configured on both units in the failover pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the failover link itself is required for replication*.

- If you use a redundant or EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal.
- You can monitor redundant or EtherChannel interfaces for failover. When an active member interface fails over to a standby interface, this activity does not cause the redundant or EtherChannel interface to appear to be failed when being monitored for device-level failover. Only when all physical interfaces fail does the redundant or EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a failover or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link. To alter the configuration, you need to either shut down the EtherChannel while you make changes, or temporarily disable failover; either action prevents failover from occurring for the duration.

Model Support

- EtherChannels are supported on ASA appliances only; they are not supported on the ASAv or the ASASM.
- Redundant interfaces are not supported on the ASASM.

Clustering

- To configure a spanned EtherChannel or an individual cluster interface, see the clustering chapter.

Redundant Interfaces

- 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.

EtherChannels

- EtherChannels are not supported on the ASAv.
- You can configure up to 48 EtherChannels.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, 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 or Cisco Nexus 7000 switch.

- The ASA does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the ASA will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch. In multiple context mode, these messages are not included in a packet capture, so that you cannot diagnose the issue easily.
- 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.

Default Settings for EtherChannels and Redundant Interfaces

This section lists default settings for interfaces if you do not have a factory default configuration.

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.
- EtherChannel port-channel interfaces—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.

Configure 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.

- [Configure a Redundant Interface, page 12-7](#)
- [Change the Active Interface, page 12-8](#)

Configure a Redundant Interface

This section describes how to create a redundant interface. By default, redundant interfaces are enabled.

Before You Begin

- 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.
- Both member interfaces must be of the same physical type. For example, both must be GigabitEthernet.
- You cannot add a physical interface to the redundant interface if you configured a name for it. You must first remove the name in the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.



Caution

If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

- Step 1** Depending on your context mode:
- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.
- Step 2** Choose **Add > Redundant Interface**.
The **Add Redundant Interface** dialog box appears.
-  **Note** In single mode, this procedure only covers a subset of the parameters on the Edit Redundant Interface dialog box; to configure other parameters, see [Chapter 15, “Routed and Transparent Mode Interfaces.”](#) Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See [Configure Multiple Contexts, page 8-15](#).
- Step 3** In the **Redundant ID** field, enter an integer between 1 and 8.
- Step 4** From the **Primary Interface** drop-down list, choose the physical interface you want to be primary.
Be sure to pick an interface that does not have a subinterface and that has not already been allocated to a context. Redundant interfaces do not support Management *slot/port* interfaces as members.
- Step 5** From the **Secondary Interface** drop-down list, choose the physical interface you want to be secondary.
- Step 6** If the interface is not already enabled, check the **Enable Interface** check box.
The interface is enabled by default.
- Step 7** To add a description, enter text in the **Description** field.

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

Step 8 Click **OK**.

You return to the **Interfaces** pane. The member interfaces now show a lock to the left of the interface ID showing that only basic parameters can be configured for it. The redundant interface is added to the table.

 GigabitEthernet0/2	Enabled	No	Redundant8	Hardware	native
GigabitEthernet0/3	Enabled	No		Hardware	native
GigabitEthernet0/3.10	Enabled	No		Logical	vlan100
GigabitEthernet0/3.11	Enabled	No		Logical	vlan11
Management0/0	Enabled	No		Hardware	native
Redundant8	Enabled	Yes		Logical	native

254710

Change the Active Interface

By default, the active interface is the first interface listed in the configuration, if it is available.

Procedure

Step 1 To view which interface is active, enter the following command in the **Tools > Command Line Interface** tool:

```
show interface redundant $\textit{number}$  detail | grep Member
```

Example:

```
show interface redundant1 detail | grep Member
Members GigabitEthernet0/3 (Active), GigabitEthernet0/2
```

Step 2 Change the active interface:

```
redundant-interface redundant $\textit{number}$  active-member  $\textit{physical\_interface}$ 
```

The **redundant $\textit{number}$** argument is the redundant interface ID, such as **redundant1**.

The **physical $\textit{interface}$** is the member interface ID that you want to be active.

Configure an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

- [Add Interfaces to the EtherChannel, page 12-9](#)
- [Customize the EtherChannel, page 12-10](#)

Add 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.

Before You Begin

- You can configure up to 48 EtherChannels.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, 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.
- To configure a spanned EtherChannel for clustering, see the clustering chapter instead of this procedure.
- All interfaces in the channel group must be the same type, speed, and duplex. Half duplex is not supported.
- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name in the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.



Caution

If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

- Step 1** Depending on your context mode:
- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.
- Step 2** Choose **Add > EtherChannel Interface**.
- The **Add EtherChannel Interface** dialog box appears.



Note

In single mode, this procedure only covers a subset of the parameters on the Edit EtherChannel Interface dialog box; to configure other parameters, see [Chapter 15, “Routed and Transparent Mode Interfaces.”](#) Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See [Configure Multiple Contexts, page 8-15](#).

- Step 3** In the **Port Channel ID** field, enter a number between 1 and 48.
- Step 4** In the *Available Physical Interface* area, click an interface and then click **Add >>** to move it to the **Members in Group** area.
- In transparent mode, if you create a channel group with multiple Management interfaces, then you can use this EtherChannel as the management-only interface.

**Note**

If you want to set the EtherChannel mode to **On**, then you must include only one interface initially. After you complete this procedure, edit the member interface, and set the mode to **On**. Apply your changes, then edit the EtherChannel to add more member interfaces.

Step 5 Repeat for each interface you want to add to the channel group.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

Step 6 Click **OK**.

You return to the **Interfaces** pane. The member interfaces now show a lock to the left of the interface ID showing that only basic parameters can be configured for it. The EtherChannel interface is added to the table.

 GigabitEthernet0/3		Disabled			Port-channel1	Hardware
Management0/0		Disabled				Hardware
Port-channel1		Enabled				EtherChannel

Step 7 Click **Apply**. All member interfaces are enabled automatically.

Related Topics

- [Link Aggregation Control Protocol, page 12-3](#)
- [Customize the EtherChannel, page 12-10](#)

Customize 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.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

Step 2 Click the port-channel interface you want to customize, and click **Edit**.

The **Edit Interface** dialog box appears.

Step 3 To override the media type, duplex, speed, and pause frames for flow control for all member interfaces, click **Configure Hardware Properties**. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

Step 4 To customize the EtherChannel, click the **Advanced** tab.

- In the **EtherChannel** area, from the **Minimum** drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 16. The default is 1.

- b. From the **Maximum** drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
- c. From the **Load Balance** drop-down list, select the criteria used to load balance the packets across the group channel interfaces. By default, the ASA balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see [Load Balancing, page 12-4](#).

Step 5 Click **OK**.

You return to the **Interfaces** pane.

Step 6 To set the mode and priority for a physical interface in the channel group:

- a. Click the physical interface in the **Interfaces** table, and click **Edit**.
The **Edit Interface** dialog box appears.
- b. Click the **Advanced** tab.
- c. In the **EtherChannel** area, from the **Mode** drop down list, choose **Active**, **Passive**, or **On**. We recommend using Active mode (the default).
- d. In the **LACP Port Priority** field, set the port priority between 1 and 65535. The default is 32768. The higher the number, the lower the priority. The ASA uses this setting to decide which interfaces are active and which are standby if you assign more interfaces than can be used. If the port priority setting is the same for all interfaces, then the priority is determined by the interface ID (slot/port). The lowest interface ID is the highest priority. For example, GigabitEthernet 0/0 is a higher priority than GigabitEthernet 0/1.

If you want to prioritize an interface to be active even though it has a higher interface ID, then set this command to have a lower value. For example, to make GigabitEthernet 1/3 active before GigabitEthernet 0/7, then make the priority value be 12345 on the 1/3 interface vs. the default 32768 on the 0/7 interface.

If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See [Step 9](#) to set the system priority.

Step 7 Click **OK**.

You return to the **Interfaces** pane.

Step 8 Click **Apply**.

Step 9 To set the LACP system priority, perform the following steps. If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See [Step 6d](#) for more information.

- a. Depending on your context mode:
 - For single mode, choose the **Configuration > Device Setup > EtherChannel** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > EtherChannel** pane.
- b. In the **LACP System Priority** field, enter a priority between 1 and 65535.
The default is 32768.

Related Topics

- [Load Balancing, page 12-4](#)
- [Add Interfaces to the EtherChannel, page 12-9](#)

Examples for EtherChannel and Redundant Interfaces

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
```

History for EtherChannels and Redundant Interfaces

Table 12-1 History for EtherChannels and Redundant Interfaces

Feature Name	Releases	Feature Information
Redundant interfaces	8.0(2)	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. 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. You can configure up to eight redundant interface pairs.

Table 12-1 *History for EtherChannels and Redundant Interfaces (continued)*

Feature Name	Releases	Feature Information
EtherChannel support	8.4(1)	You can configure up to 48 802.3ad EtherChannels of eight active interfaces each. We modified or introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface Configuration > Device Setup > EtherChannel Note EtherChannel is not supported on the ASA 5505.
Support for 16 active links in an EtherChannel	9.2(1)	You can now configure up to 16 active links in an EtherChannel. Previously, you could have 8 active links and 8 standby links. Be sure that your switch can support 16 active links (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module). Note If you upgrade from an earlier ASA version, the maximum active interfaces is set to 8 for compatibility purposes. We modified the following screen: Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface > Advanced.



VLAN Interfaces

This chapter tells how to configure VLAN subinterfaces.



Note

For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.

- [About VLAN Interfaces, page 13-1](#)
- [Licensing for VLAN Interfaces, page 13-1](#)
- [Guidelines for VLAN Interfaces, page 13-2](#)
- [Default Settings for VLAN Interfaces, page 13-2](#)
- [Configure VLAN Subinterfaces and 802.1Q Trunking, page 13-3](#)
- [Examples for VLAN Interfaces, page 13-4](#)
- [History for VLAN Interfaces, page 13-4](#)

About VLAN Interfaces

VLAN 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.

Licensing for VLAN Interfaces

Model	License Requirement
ASA 5506-X series	Base License: 5
	Security Plus License: 30
ASA 5508-X	Base License: 50

Model	License Requirement
ASA 5512-X	Base License: 50 Security Plus License: 100
ASA 5515-X	Base License: 100
ASA 5516-X	Base License: 50
ASA 5525-X	Base License: 200
ASA 5545-X	Base License: 300
ASA 5555-X	Base License: 500
ASA 5585-X	Base and Security Plus License: 1024
ASAv5 and ASAv10	Standard and Premium License: 50
ASAv30	Standard and Premium License: 200

**Note**

For an interface to count against the VLAN limit, you must assign a VLAN to it.

Guidelines for VLAN Interfaces

Model Support

VLAN subinterfaces are not supported on the ASASM; ASASM interfaces are already VLAN interfaces assigned from the switch.

Additional Guidelines

- 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 and for EtherChannel links. Because the physical, redundant, or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical, redundant, or EtherChannel interface does not pass traffic by not configuring a name for the interface. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the name as usual.
- (All models except for the ASA 5585-X) You cannot configure subinterfaces on the Management interface.
- The ASA does not support the Dynamic Trunking Protocol (DTP), so you must configure the connected switch port to trunk unconditionally.

Default Settings for VLAN Interfaces

This section lists default settings for interfaces if you do not have a factory default configuration.

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.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.

Configure VLAN Subinterfaces and 802.1Q Trunking

Add a VLAN subinterface to a physical, redundant, or EtherChannel interface.

Before You Begin

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

Procedure

-
- Step 1** Depending on your context mode:
- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
 - For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.
- Step 2** Choose **Add > Interface**.
- The **Add Interface** dialog box appears.
- 
- Note** In single mode, this procedure only covers a subset of the parameters on the Edit Interface dialog box; to configure other parameters, see [Chapter 15, “Routed and Transparent Mode Interfaces.”](#) Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See [Configure Multiple Contexts, page 8-15](#).
-
- Step 3** From the **Hardware Port** drop-down list, choose the physical, redundant, or port-channel interface to which you want to add the subinterface.
- Step 4** If the interface is not already enabled, check the **Enable Interface** check box.
- The interface is enabled by default.
- Step 5** In the **VLAN ID** field, enter the VLAN ID between 1 and 4095.
- Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information. For multiple context mode, you can only set the VLAN in the system configuration.
- Step 6** In the **Subinterface ID** field, enter the subinterface ID as an integer between 1 and 4294967293.

The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

Step 7 (Optional) In the **Description** field, enter a description for this interface.

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

Step 8 Click **OK**.

You return to the **Interfaces** pane.

Related Topics

- [Licensing for VLAN Interfaces, page 13-1](#)

Examples for VLAN Interfaces

The following example configures parameters for a subinterface in single mode:

```
interface gigabitethernet 0/1.1
  vlan 101
  no shutdown
```

History for VLAN Interfaces

Table 13-1 History for VLAN Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license VLANs from 0 to 10. • ASA5510 Security Plus license VLANs from 10 to 25. • ASA5520 VLANs from 25 to 100. • ASA5540 VLANs from 100 to 200.
Increased VLANs	7.2(2)	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).
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.



VXLAN Interfaces

This chapter tells how to configure Virtual eXtensible LAN (VXLAN) interfaces. VXLANs act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks.

- [About VXLAN Interfaces, page 14-1](#)
- [Guidelines for VXLAN Interfaces, page 14-6](#)
- [Default Settings for VXLAN Interfaces, page 14-6](#)
- [Configure VXLAN Interfaces, page 14-6](#)
- [Examples for VXLAN Interfaces, page 14-8](#)
- [History for VXLAN Interfaces, page 14-12](#)

About VXLAN Interfaces

VXLAN provides the same Ethernet Layer 2 network services as VLAN does, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- Flexible placement of multitenant segments throughout the data center.
- Higher scalability to address more Layer 2 segments: up to 16 million VXLAN segments.

This section describes how VXLAN works. For detailed information, see RFC 7348.

- [VXLAN Encapsulation, page 14-1](#)
- [VXLAN Tunnel Endpoint, page 14-2](#)
- [VTEP Source Interface, page 14-2](#)
- [VNI Interfaces, page 14-2](#)
- [Peer VTEPs, page 14-3](#)
- [VXLAN Use Cases, page 14-3](#)


VXLAN Encapsulation

VXLAN is a Layer 2 overlay scheme on a Layer 3 network. It uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. The original Layer 2 frame has a VXLAN header added and is then placed in a UDP-IP packet.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces to which you apply your security policy, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

The following figure shows two ASAs and Virtual Server 2 acting as VTEPs across a Layer 3 network, extending the VNI 1, 2, and 3 networks between sites. The ASAs act as bridges or gateways between VXLAN and non-VXLAN networks.



The underlying IP network between VTEPs is independent of the VXLAN overlay. Encapsulated packets are routed based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address. The destination IP address can be a multicast group when the remote VTEP is not known. The destination port is UDP port 4789 by default (user configurable).

VTEP Source Interface

The VTEP source interface is a regular ASA interface (physical, redundant, EtherChannel, or even VLAN) with which you plan to associate all VNI interfaces. You can configure one VTEP source interface per ASA/security context.

The VTEP source interface can be devoted wholly to VXLAN traffic, although it is not restricted to that use. If desired, you can use the interface for regular traffic and apply a security policy to the interface for that traffic. For VXLAN traffic, however, all security policy must be applied to the VNI interfaces. The VTEP interface serves as a physical port only.

In transparent firewall mode, the VTEP source interface is not part of a BVI, and you do configure an IP address for it, similar to the way the management interface is treated.

VNI Interfaces

VNI interfaces are similar to VLAN interfaces: they are virtual interfaces that keep network traffic separated on a given physical interface by using tagging. You apply your security policy directly to each VNI interface.

All VNI interfaces are associated with the same VTEP interface.

VXLAN Packet Processing

Traffic entering and exiting the VTEP source interface is subject to VXLAN processing, specifically encapsulation or decapsulation.

Encapsulation processing includes the following tasks:

- The VTEP source interface encapsulates the inner MAC frame with the VXLAN header.
- The UDP checksum field is set to zero.
- The Outer frame source IP is set to the VTEP interface IP.
- The Outer frame destination IP is decided by a remote VTEP IP lookup.

Decapsulation; the ASA only decapsulates a VXLAN packet if:

- It is a UDP packet with the destination port set to 4789 (this value is user configurable).
- The ingress interface is the VTEP source interface.
- The ingress interface IP address is the same as the destination IP address.
- The VXLAN packet format is compliant with the standard.

Peer VTEPs

When the ASA sends a packet to a device behind a peer VTEP, the ASA needs two important pieces of information:

- The destination MAC address of the remote device
- The destination IP address of the peer VTEP

There are two ways in which the ASA can find this information:

- A single peer VTEP IP address can be statically configured on the ASA.

You cannot manually define multiple peers.

The ASA then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC address.

- A multicast group can be configured on each VNI interface (or on the VTEP as a whole).

The ASA sends a VXLAN-encapsulated ARP broadcast packet within an IP multicast packet through the VTEP source interface. The response to this ARP request enables the ASA to learn both the remote VTEP IP address along with the destination MAC address of the remote end node.

The ASA maintains a mapping of destination MAC addresses to remote VTEP IP addresses for the VNI interfaces.

VXLAN Use Cases


This section describes the use cases for implementing VXLAN on the ASA.

- [VXLAN Bridge or Gateway Overview, page 14-4](#)
- [VXLAN Bridge \(Transparent Mode\), page 14-4](#)
- [VXLAN Gateway \(Routed Mode\), page 14-4](#)
- [Router Between VXLAN Domains, page 14-5](#)

VXLAN Bridge or Gateway Overview


Each ASA VTEP acts as a bridge or gateway between end nodes such as VMs, servers, and PCs and the VXLAN overlay network. For incoming frames received with VXLAN encapsulation over the VTEP source interface, the ASA strips out the VXLAN header and forwards it to a physical interface connected to a non-VXLAN network based on the destination MAC address of the inner Ethernet frame.

The ASA always processes VXLAN packets; it does not just forward VXLAN packets untouched between two other VTEPs.




VXLAN Bridge (Transparent Mode)

In transparent firewall mode, the ASA can serve as a VXLAN bridge between a (remote) VXLAN segment and a local segment where both are in the same network. In this case, one member of the bridge virtual interface (BVI) is a regular interface while the other member is a VNI interface.




VXLAN Gateway (Routed Mode)

The ASA can serve as a router between VXLAN and non-VXLAN domains, connecting devices on different networks.



Router Between VXLAN Domains

With a VXLAN-stretched Layer 2 domain, a VM can point to an ASA as its gateway while the ASA is not on the same rack, or even when the ASA is far away over the Layer 3 network.



See the following notes about this scenario:

1. For packets from VM3 to VM1, the destination MAC address is the ASA MAC address, because the ASA is the default gateway.
2. The VTEP source interface on Virtual Server 2 receives packets from VM3, then encapsulates the packets with VNI 3's VXLAN tag and sends them to the ASA.
3. When the ASA receives the packets, it decapsulates the packets to get the inner frames.
4. The ASA uses the inner frames for route lookup, then finds that the destination is on VNI 2. If it does not already have a mapping for VM1, the ASA sends an encapsulated ARP broadcast on the multicast group IP on VNI 2.



Note The ASA must use dynamic VTEP peer discovery because it has multiple VTEP peers in this scenario.

5. The ASA encapsulates the packets again with the VXLAN tag for VNI 2 and sends the packets to Virtual Server 1. Before encapsulation, the ASA changes the inner frame destination MAC address to be the MAC of VM1 (multicast-encapsulated ARP might be needed for the ASA to learn the VM1 MAC address).
6. When Virtual Server 1 receives the VXLAN packets, it decapsulates the packets and delivers the inner frames to VM1.

Guidelines for VXLAN Interfaces

IPv6

- The VNI interface supports IPv6 traffic, but the VTEP source interface IP address only supports IPv4.
- IPv6 OSPF interface settings are not supported.

Clustering

ASA clustering does not support VXLAN in Individual Interface mode. Only Spanned EtherChannel mode supports VXLAN.

Routing

- Only static routing is supported on the VNI interface; dynamic routing protocols are not supported.
- Policy Based Routing is not supported.

Default Settings for VXLAN Interfaces

VNI interfaces are enabled by default.

Configure VXLAN Interfaces

To configure VXLAN, perform the following steps.

-
- Step 1** [Configure the VTEP Source Interface, page 14-7.](#)
- Step 2** [Configure the VNI Interface, page 14-8](#)
-

Configure the VTEP Source Interface

You can configure one VTEP source interface per ASA or per security context. The VTEP is defined as a Network Virtualization Endpoint (NVE); VXLAN VTEP is the only supported NVE at this time.

Before You Begin

For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**, and edit the interface you want to use for the VTEP source interface.
- Step 2** (Transparent Mode) Check the **VTEP Source Interface** check box.
- This setting lets you configure an IP address for the interface. This command is optional for routed mode where this setting restricts traffic to VXLAN only on this interface.
- Step 3** Configure the source interface name and IPv4 address, and click **OK**.
- Step 4** Choose **Configuration > Device Setup > Interface Settings > VXLAN**.
- Step 5** (Optional) Enter a **VXLAN Destination Port** value if you want to change from the default 4789.
- In multiple context mode, configure this setting in the System execution space.
- Step 6** Check the **Enable Network Virtualization Endpoint encapsulation using VXLAN** check box.
- Step 7** Choose the **VTEP Tunnel Interface** from the drop-down list.
- Step 8** (Optional) Check the **Configure Packet Recipient** check box.
- (Multiple context mode; Optional for single mode) Enter the **Specify Peer VTEP IP Address** to manually specify the peer VTEP IP address
- If you specify the peer IP address, you cannot use multicast group discovery. Multicast is not supported in multiple context mode, so manual configuration is the only option. You can only specify one peer for the VTEP.
- (Single mode only) Enter the **Multicast traffic to default multicast address** to specify a default multicast group for all associated VNI interfaces.
- If you do not configure the multicast group per VNI interface, then this group is used. If you configure a group at the VNI interface level, then that group overrides this setting.
- Step 9** Click **Apply**.
-

Configure the VNI Interface

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.


-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**, and click **Add > VNI Interface**.
- Step 2** Enter the **VNI ID**, between 1 and 10000.
This ID is just an internal interface identifier.
- Step 3** Enter the **VNI Segment ID**, between 1 and 16777215.
The segment ID is used for VXLAN tagging.
- Step 4** (Transparent Mode) Choose the **Bridge Group** to which you want to assign this interface.
See [Configure Transparent Mode Interfaces, page 15-9](#) to configure the BVI interface and associate regular interfaces to this bridge group.
- Step 5** Enter the **Interface Name**.
The name 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.
- Step 6** Enter the **Security Level**, between 0 (lowest) and 100 (highest). See [Security Levels, page 15-2](#).
- Step 7** (Single Mode) Enter the **Multicast Group IP Address**.
If you do not set the multicast group for the VNI interface, the default group from the VTEP source interface configuration is used, if available. If you manually set a VTEP peer IP for the VTEP source interface, you cannot specify a multicast group for the VNI interface. Multicast is not supported in multiple context mode.
- Step 8** Check the **NVE Mapped to VTEP Interface** check box.
This setting associates the VNI interface with the VTEP source interface.
- Step 9** Check the **Enable Interface** check box. This setting is enabled by default.
- Step 10** (Routed Mode) In the **IP Address** area, configure an IPv4 address. To configure IPv6, click the **IPv6** tab.
- Step 11** Click **OK**, and then **Apply**.
-

Examples for VXLAN Interfaces

See the following configuration examples for VXLAN.

- [Transparent VXLAN Gateway Example, page 14-9](#)
- [VXLAN Routing Example, page 14-11](#)

Transparent VXLAN Gateway Example



See the following description of this example:

- The outside interface on GigabitEthernet 0/0 is used as the VTEP source interface, and it is connected to the Layer 3 network.
- The insidevm100 VLAN subinterface on GigabitEthernet 0/1.100 is connected to the 10.10.10.0/24 network, on which VM3 resides. When VM3 communicates with VM1 (not shown; both have 10.10.10.0/24 IP addresses), the ASA uses VXLAN tag 6000.
- The insidevm200 VLAN subinterface on GigabitEthernet 0/1.200 is connected to the 10.20.20.0/24 network, on which VM2 resides. When VM2 communicates with VM4 (not shown; both have 10.20.20.0/24 IP addresses), the ASA uses VXLAN tag 8000.
- The insidepc interface on GigabitEthernet 0/2 is connected to the 10.30.30.0/24 network on which a few PCs reside. When those PCs communicate with VMs/PCs (not shown) behind a remote VTEP that belongs to same network (all have 10.30.30.0/24 IP addresses), the ASA uses VXLAN tag 10000.

ASA Configuration

```
firewall transparent
vxlan port 8427
!
interface gigabitethernet0/0
  nve-only
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
nve 1
  encapsulation vxlan
  source-interface outside
!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  bridge-group 1
  vtep-nve 1
```

```


    mcast-group 235.0.0.100
!
interface vni2
    segment-id 8000
    nameif vxlan8000
    security-level 0
    bridge-group 2
    vtep-nve 1
    mcast-group 236.0.0.100
!
interface vni3
    segment-id 10000
    nameif vxlan10000
    security-level 0
    bridge-group 3
    vtep-nve 1
    mcast-group 236.0.0.100
!
interface gigabitethernet0/1.100
    nameif insidevm100
    security-level 100
    bridge-group 1
!
interface gigabitethernet0/1.200
    nameif insidevm200
    security-level 100
    bridge-group 2
!
interface gigabitethernet0/2
    nameif insidepc
    security-level 100
    bridge-group 3
!
interface bvi 1
    ip address 10.10.10.1 255.255.255.0
!
interface bvi 2
    ip address 10.20.20.1 255.255.255.0
!
interface bvi 3
    ip address 10.30.30.1 255.255.255.0

```

Notes

- For VNI interfaces vni1 and vni2, the inner VLAN tag is removed during encapsulation.
- VNI interfaces vni2 and vni3 share the same multicast IP address for encapsulated ARP over multicast. This sharing is allowed.
- The ASA bridges the VXLAN traffic to non-VXLAN-supported interfaces based on the above BVIs and bridge group configurations. For each of the stretched Layer 2 network segments (10.10.10.0/24, 10.20.20.0/24 and 10.30.30.0/24), the ASA serves as a bridge.
- It is allowed to have more than one VNI or more than one regular interface (VLAN or just physical interface) in a bridge group. The forwarding or association between VXLAN segment ID to the VLAN ID (or a physical interface) is decided by the destination MAC address and which interface connects to the destination.
- The VTEP source-interface is a Layer 3 interface in transparent firewall mode indicated by **nve-only** in the interface configuration. The VTEP source interface is not a BVI interface or a management interface, but it has an IP address and uses the routing table.

VXLAN Routing Example



See the following description of this example:

- VM1 (10.10.10.10) is hosted on Virtual Server 1, and VM2 (10.20.20.20) is hosted on Virtual Server 2.
- The default gateway for VM1 is the ASA, which is not in the same pod as Virtual Server 1, but VM1 is not aware of it. VM1 only knows that its default gateway IP address is 10.10.10.1. Similarly, VM2 only knows that its default gateway IP address is 10.20.20.1.
- The VTEP-supported hypervisors on Virtual Server 1 and 2 are able to communicate with the ASA over the same subnet or through a Layer 3 network (not shown; in which case, the ASA and uplinks of virtual servers have different network addresses).
- VM1's packet will be encapsulated by its hypervisor's VTEP and sent to its default gateway over VXLAN tunneling.
- When VM1 sends a packet to VM2, the packet will be sent through default gateway 10.10.10.1 from its perspective. Virtual Server1 knows 10.10.10.1 is not local, so the VTEP encapsulates the packet over VXLAN and sends it to ASA's VTEP.
- On the ASA, the packet is decapsulated. The VXLAN segment ID is learned during decapsulation. The ASA then re-injects the inner frame to the corresponding VNI interface (vni1) based on the VXLAN segment ID. The ASA then conducts a route lookup and sends the inner packet through another VNI interface, vni2. All egressing packets through vni2 are encapsulated with VXLAN segment 8000 and sent through the VTEP to outside.
- Eventually the encapsulated packet is received by the VTEP of Virtual Server 2, which decapsulates it and forwards it to VM2.

ASA Configuration

```
interface gigabitethernet0/0
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
vxe 1
  encapsulation vxlan
  source-interface outside
  default-mcast-group 235.0.0.100
```

```

!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  vtep-nve 1
  ip address 10.20.20.1 255.255.255.0
!
interface vni2
  segment-id 8000
  nameif vxlan8000
  security-level 0
  vtep-nve 1
  ip address 10.10.10.1 255.255.255.0
!

```

History for VXLAN Interfaces

Table 14-1 History for VXLAN Interfaces

Feature Name	Releases	Feature Information
VXLAN support	9.4(1)	VXLAN support was added, including VXLAN tunnel endpoint (VTEP) support. You can define one VTEP source interface per ASA or security context. We introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces > Add > VNI Interface Configuration > Device Setup > Interface Settings > VXLAN



Routed and Transparent Mode Interfaces

This chapter includes tasks to complete the interface configuration for all models in routed or transparent firewall mode.

- [About Routed and Transparent Mode Interfaces, page 15-1](#)
- [Guidelines for Routed and Transparent Mode Interfaces, page 15-4](#)
- [Defaults for Routed and Transparent Mode Interfaces, page 15-5](#)
- [Configure Routed Mode Interfaces, page 15-6](#)
- [Configure Transparent Mode Interfaces, page 15-9](#)
- [Configure IPv6 Addressing, page 15-12](#)
- [Monitoring Routed and Transparent Mode Interfaces, page 15-15](#)
- [Examples for Routed and Transparent Mode Interfaces, page 15-16](#)
- [History for Routed and Transparent Mode Interfaces, page 15-17](#)



Note

For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

About Routed and Transparent Mode Interfaces

When the ASA is in routed firewall mode (the default), each interface is a Layer 3 routed interface for which you need to set an IP address on a unique subnet. All transparent mode interfaces belong to bridge groups.

- [Security Levels, page 15-2](#)
- [Bridge Groups in Transparent Mode, page 15-2](#)
- [PPPoE \(Routed Mode Only\), page 15-3](#)
- [Dual IP Stack \(IPv4 and IPv6\), page 15-3](#)
- [IPv6, page 15-3](#)

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.

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, 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.

Related Topics

[Allow Same Security Level Communication, page 16-7](#)

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 Cisco 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.

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 [Management Interface, page 11-2](#) for more information.

PPPoE (Routed Mode Only)

PPPoE combines two widely accepted standards, Ethernet and PPP, to provide an authenticated method of assigning IP addresses to client systems. PPPoE clients are typically personal computers connected to an ISP over a remote broadband connection, such as DSL or cable service. ISPs deploy PPPoE because it supports high-speed broadband access using their existing remote access infrastructure and because it is easier for customers to use.

PPPoE provides a standard method of employing the authentication methods of the Point-to-Point Protocol (PPP) over an Ethernet network. When used by ISPs, PPPoE allows authenticated assignment of IP addresses. In this type of implementation, the PPPoE client and server are interconnected by Layer 2 bridging protocols running over a DSL or other broadband connection.

PPPoE is composed of two main phases:

- **Active Discovery Phase**—In this phase, the PPPoE client locates a PPPoE server, called an access concentrator. During this phase, a Session ID is assigned and the PPPoE layer is established.
- **PPP Session Phase**—In this phase, PPP options are negotiated and authentication is performed. Once the link setup is completed, PPPoE functions as a Layer 2 encapsulation method, allowing data to be transferred over the PPP link within PPPoE headers.

At system initialization, the PPPoE client establishes a session with the access concentrator by exchanging a series of packets. Once the session is established, a PPP link is set up, which includes authentication using Password Authentication protocol (PAP). Once the PPP session is established, each packet is encapsulated in the PPPoE and PPP headers.

Dual IP Stack (IPv4 and IPv6)

The Cisco 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.

IPv6

This section includes information about how to configure IPv6.

- [IPv6 Addressing, page 15-3](#)
- [Modified EUI-64 Interface IDs, page 15-4](#)
- [Unsupported IPv6 Commands in Transparent Mode, page 15-4](#)

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. For transparent mode, 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. In transparent mode, 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 the interface, so you do not also need to specifically configure a link-local address. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Modified EUI-64 Interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The ASA can enforce this requirement for hosts attached to the local link.

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
%ASA-3-325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address.

Unsupported IPv6 Commands in Transparent Mode

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**

Guidelines for Routed and Transparent Mode Interfaces

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 [Configure Multiple Contexts, page 8-15](#).
- PPPoE is not supported in multiple context mode.
- For multiple context mode in transparent mode, each context must use different interfaces; you cannot share an interface across contexts.

- For multiple context mode in transparent 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

Do not configure failover interfaces with the procedures in this chapter. See [Chapter 7, “Failover for High Availability,”](#) to configure the failover and state links.

IPv6

The ASA does not support IPv6 anycast addresses.

Model Support

PPPoE and DHCP are not supported on the ASASM.

VLAN IDs 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.

Transparent Mode Guidelines

- You can configure up to 250 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.
- 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.
- 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.
- PPPoE is not supported for the management interface in transparent mode.

Defaults for Routed and Transparent Mode Interfaces

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.

Configure Routed Mode Interfaces

To configure routed mode interfaces, perform the following steps.

-
- Step 1** [Configure General Routed Mode Interface Parameters, page 15-6.](#)
 - Step 2** [Configure PPPoE, page 15-8.](#)
 - Step 3** [Configure IPv6 Addressing, page 15-12.](#)
-

Configure General Routed Mode Interface Parameters

This procedure describes how to set the name, security level, IPv4 address, and other options.

Before You Begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
 - Step 2** Choose the interface row, and click **Edit**.
The **Edit Interface** dialog box appears with the **General** tab selected.
 - Step 3** In the **Interface Name** field, enter a name up to 48 characters in length.
 - Step 4** In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).
 - Step 5** (Optional; not supported for redundant interfaces) To set this interface as a management-only interface, check the **Dedicate this interface to management-only** check box.

Through traffic is not accepted on a management-only interface.

(All ASAs except for the ASA 5585-X) You cannot disable this option on the Management interface.



Note The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

- Step 6** If the interface is not already enabled, check the **Enable Interface** check box.
- Step 7** To set the IP address, use one of the following options.



Note For use with failover, you must set the IP address and standby address manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.

- a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned.

- b. To use a generated string for option 61, click **Use “Cisco-<MAC>-<interface_name>-<host>”**.
- c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
- d. (Optional) To assign an administrative distance to the learned route, enter a value between 1 and 255 in the **DHCP Learned Route Metric** field. If this field is left blank, the administrative distance for the learned routes is 1.
- e. (Optional) To enable tracking for DHCP-learned routes, check **Enable Tracking for DHCP Learned Routes**. Set the following values:

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

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



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

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

Monitor Options—Click this button to open the **Route Monitoring Options** dialog box. In the **Route Monitoring Options** dialog box you can configure the parameters of the tracked object monitoring process.

- f. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

- g. (Optional) To renew the lease, click **Renew DHCP Lease**.
- (Single mode only) To obtain an IP address using PPPoE, check **Use PPPoE**.
 - a. In the **Group Name** field, specify a group name.
 - b. In the **PPPoE Username** field, specify the username provided by your ISP.
 - c. In the **PPPoE Password** field, specify the password provided by your ISP.
 - d. In the **Confirm Password** field, retype the password.
 - e. For PPP authentication, click either the **PAP**, **CHAP**, or **MSCHAP** radio button.

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

- f. (Optional) To store the username and password in flash memory, check the **Store Username and Password in Local Flash** check box.

The ASA stores the username and password in a special location of NVRAM. If an Auto Update Server sends a **clear configure** command to the ASA, and the connection is then interrupted, the ASA can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

- g. (Optional) To display the **PPPoE IP Address and Route Settings** dialog box where you can choose addressing and tracking options, click **IP Address and Route Settings**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Related Topics

- [Configure the MAC Address, MTU, and TCP MSS, page 16-5](#)
- [Configure IPv6 Addressing, page 15-12](#)
- [Enable the Physical Interface and Configure Ethernet Parameters, page 11-6](#)
- [Configure PPPoE, page 15-8](#)

Configure PPPoE

If you use PPPoE for IP addressing, you can manually configure the IP address and configure route settings.

Procedure

Step 1 Choose **Configuration > Interfaces > Add/Edit Interface > General**, and then click **PPPoE IP Address and Route Settings**.

Step 2 In the **IP Address** area, choose one of the following:

- **Obtain IP Address using PPP**—Dynamically configure the IP address.
- **Specify an IP Address**—Manually configure the IP address.

Step 3 In the **Route Settings Area**, configure the following:

- **Obtain default route using PPPoE**—Set the default routes when the PPPoE client has not yet established a connection. When using this option, you cannot have a statically defined route in the configuration.
- **PPPoE learned route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrative distance for the learned routes is 1.
- **Enable tracking**—Enable route tracking for PPPoE-learned routes. Route tracking is only available in single, routed mode.
- **Primary Track**—Configure the primary PPPoE route tracking.

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

Step 4 Click **OK**.

Configure Transparent Mode Interfaces

To configure bridge groups and associated interfaces, perform these steps.

-
- Step 1** [Configure Bridge Groups, page 15-9.](#)
- Step 2** [Configure General Transparent Mode Interface Parameters, page 15-10.](#)
- Step 3** [Configure a Management Interface, page 15-11.](#)
- Step 4** [Configure IPv6 Addressing, page 15-12.](#)
-

Configure 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.



Note

For a separate management interface (for supported models), a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Procedure

- Step 1** Choose **Configuration > Interfaces**, and then choose **Add > Bridge Group Interface**.
- Step 2** In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.
- Step 3** In the **IP Address** field, enter the management IPv4 address.

- Step 4** In the Subnet Mask field, enter the subnet mask or choose one from the menu.
- Do not assign a host address (/32 or 255.255.255.255) to the transparent firewall. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The ASA drops all ARP packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the ASA drops the ARP request from the downstream router to the upstream router.
- Step 5** (Optional) In the **Description** field, enter a description for this bridge group.
- Step 6** Click **OK**.
- Step 7** A Bridge Group Virtual Interface (BVI) is added to the interface table, along with the physical and subinterfaces.
-

Configure General Transparent Mode Interface Parameters

This procedure describes how to set the name, security level, and bridge group for each transparent interface.

Before You Begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.
- Do not use this procedure for Management interfaces; see [Configure a Management Interface, page 15-11](#) to configure the Management interface.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- BVIs appear in the table alongside physical interfaces, subinterfaces, redundant interfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.
- Step 2** Choose the row for a non-BVI interface, and click **Edit**.
- The **Edit Interface** dialog box appears with the **General** tab selected.
- Step 3** In the **Bridge Group** drop-down menu, choose the bridge group to which you want to assign this interface.
- Step 4** In the **Interface Name** field, enter a name up to 48 characters in length.
- Step 5** In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).
- Step 6** If the interface is not already enabled, check the **Enable Interface** check box.



Note The **Channel Group** field is read-only and indicates if the interface is part of an EtherChannel.

- Step 7** (Optional) If you install a module, and you want to demonstrate the module functionality on a non-production ASA, check the **Forward traffic to the ASA module for inspection and reporting check box**. See the module chapter or quick start guide for more information.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Configure a Management Interface

You can configure one management interface separate from the bridge group interfaces in single mode or per context.

Before You Begin

- 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 ASASM.)
- In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. You must connect to a data interface.
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration > Device Setup > Interface Settings > Interfaces**.

Step 2 Choose the row for a Management interface, subinterface, or EtherChannel port-channel interface comprised of Management interfaces, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 In the **Bridge Group** drop-down menu, leave the default **--None--**. You cannot assign a management interface to a bridge group.

Step 4 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 5 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).



Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

**Note**

For use with failover, you must set the IP address and standby address manually; DHCP is not supported. Set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
 - a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned.

- b. To use a generated string for option 61, click **Use “Cisco-<MAC>-<interface_name>-<host>”**.
- c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
- d. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

- e. (Optional) To renew the lease, click **Renew DHCP Lease**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns.

Step 9 Click **OK**.

Related Topics

[Management Interface, page 11-2](#)

Configure IPv6 Addressing

This section describes how to configure IPv6 addressing in routed and transparent mode.

- [Configure a Global IPv6 Address, page 15-12](#)
- (Optional) [Configure the Link-Local Addresses Automatically, page 15-14](#)
- (Optional) [Configure the Link-Local Addresses Manually, page 15-14](#)

Configure a Global IPv6 Address

To configure a global IPv6 address for any routed mode interface and for the transparent mode bridge group interface and 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.

**Note**

See [Chapter 29, “IPv6 Neighbor Discovery,”](#) to configure IPv6 neighbor discovery.

Before You Begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration > Device Setup > Interface Settings > Interfaces**.

Step 2 Choose an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

In transparent mode, select a bridge group interface or management interface only.

Step 3 Click the **IPv6** tab.

Step 4 Check the **Enable IPv6** check box.

Step 5 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

Step 6 Configure the global IPv6 address using one of the following methods.

- (Routed Mode only) Stateless autoconfiguration—In the **Interface IPv6 Addresses** area, check the **Enable address autoconfiguration** check box.

Enabling stateless autoconfiguration on the interface configures IPv6 addresses based upon prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

**Note**

Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the ASA does send Router Advertisement messages in this case. See the **Suppress RA** check box to suppress messages.

- Manual configuration—To manually configure a global IPv6 address:

- a. In the **Interface IPv6 Addresses** area, click **Add**.

The **Add IPv6 Address for Interface** dialog box appears.

- b. In the **Address/Prefix Length** field, enter either full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. (Routed Mode only) If you only enter the prefix, then be sure to check the **EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).
- c. Click **OK**.

- Step 7** (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, refer to the [Configure the IPv6 Prefix in Router Advertisements, page 29-11](#).
- Step 8** Click **OK**.
- You return to the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
-

(Optional) Configure the Link-Local Addresses Automatically

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required for the interface ID.)

To automatically configure the link-local addresses for an interface, perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select an interface, and click **Edit**.
- The **Edit Interface** dialog box appears with the **General** tab selected.
- In transparent mode, select any non-bridge group interface.
- Step 3** Click the **IPv6** tab.
- Step 4** In the **IPv6 configuration** area, check the **Enable IPv6** check box.
- This option enables IPv6 and automatically generates the link-local address using the Modified EUI-64 interface ID based on the interface MAC address.
- Step 5** Click **OK**.
-

(Optional) Configure the Link-Local Addresses Manually

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

To assign a link-local address to an interface, perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select an interface, and click **Edit**.
- The **Edit Interface** dialog box appears with the **General** tab selected.
- In transparent mode, select any non-bridge group interface.

- Step 3** Click the **IPv6** tab.
- Step 4** To set the link-local address, enter an address in the **Link-local address** field.
A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82.
See [IPv6 Addresses, page 41-5](#) for more information about IPv6 addressing.
- Step 5** Click **OK**.
-

Monitoring Routed and Transparent Mode Interfaces

You can monitor interface statistics, status, PPPoE, and more.

- [Interface Statistics, page 15-15](#)
- [Static Route Tracking, page 15-15](#)
- [PPPoE, page 15-15](#)
- [DHCP, page 15-16](#)
- [Dynamic ACLs, page 15-16](#)

Interface Statistics

- **Monitoring > Interfaces > Interface Graphs**

Lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA shows only statistics for the current context. The number of statistics shown for a subinterface is a subset of the number of statistics shown for a physical interface.

- **Monitoring > Interfaces > Interface Graphs > Graph/Table**

Shows a graph for the selected statistics. The Graph window can show up to four graphs and tables at a time. By default, the graph or table displays the real-time statistics. If you enable History Metrics, you can view statistics for past time periods.

Static Route Tracking

- **Monitoring > Interfaces > interface connection > Track Status**

Displays information about the tracked object.

- **Monitoring > Interfaces > interface connection > Monitoring Statistics**

Displays statistics for the SLA monitoring process.

PPPoE

- **Monitoring > Interfaces > PPPoE Client > PPPoE Client Lease Information**

Displays information about current PPPoE connections.

DHCP

- **Monitoring > Interfaces > DHCP > DHCP Server Table**
Lists the IP addresses assigned to DHCP clients.
- **Monitoring > Interfaces > DHCP > DHCP Server Table > DHCP Client Lease Information**
Shows information about the DHCP lease.
- **Monitoring > Interfaces > DHCP > DHCP Statistics**
Shows statistics for the DHCP server feature.

Dynamic ACLs

Monitoring > Interfaces > Dynamic ACLs

Shows a table of the Dynamic ACLs, which are functionally identical to the user-configured ACLs except that they are created, activated and deleted automatically by the ASA. These ACLs do not show up in the configuration and are only visible in this table. They are identified by the “(dynamic)” keyword in the ACL header.

Examples for Routed and Transparent Mode Interfaces

The following example for transparent mode 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

```

History for Routed and Transparent Mode Interfaces

Table 15-1 History for Interfaces in Transparent Mode

Feature Name	Platform Releases	Feature Information
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
Bridge groups for transparent mode	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We modified or introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface
Transparent mode bridge group maximum increased to 250	9.3(1)	The bridge group maximum was increased from 8 to 250 bridge groups. You can configure up to 250 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. We modified the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface



Advanced Interface Configuration

This chapter describes how to configure MAC addresses for interfaces, how to set the maximum transmission unit (MTU), and set the TCP maximum segment size (TCP MSS), and how to allow same security level communication. Setting the correct MTU and maximum TCP segment size is essential for the best network performance.

- [About Advanced Interface Configuration, page 16-1](#)
- [Configure the MAC Address, MTU, and TCP MSS, page 16-5](#)
- [Allow Same Security Level Communication, page 16-7](#)
- [Monitoring the ARP and MAC Address Table, page 16-7](#)



Note

For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

About Advanced Interface Configuration

This section describes advanced interface settings.

- [About MAC Addresses, page 16-1](#)
- [About the MTU, page 16-2](#)
- [About the TCP MSS, page 16-3](#)
- [Inter-Interface Communication, page 16-4](#)
- [Intra-Interface Communication \(Routed Firewall Mode\), page 16-4](#)

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. You can assign each MAC address manually, or you can automatically generate MAC addresses for shared interfaces in contexts. 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.

Related Topics

- [How the ASA Classifies Packets, page 8-3](#)
- [Assign MAC Addresses to Context Interfaces Automatically, page 8-21](#)

About the MTU

The MTU specifies the maximum frame payload size that the ASA can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, FCS, or VLAN tagging. The Ethernet header is 14 bytes and the FCS is 4 bytes. When you set the MTU to 1500, the expected frame size is 1518 bytes including the headers. If you are using VLAN tagging (which adds an additional 4 bytes), then when you set the MTU to 1500, the expected frame size is 1522. Do not set the MTU value higher to accommodate these headers. For information about accommodating TCP headers for encapsulation, do not alter the MTU setting; instead change the TCP Maximum Segment Size.

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.

- [Default MTU, page 16-3](#)
- [Path MTU Discovery, page 16-3](#)
- [MTU and Jumbo Frames, page 16-3](#)

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.

MTU and Jumbo Frames

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 you enable jumbo frames, you can set the MTU up to 9198 bytes.

About the TCP MSS

The 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, page 16-3](#)
- [TCP MSS for VPN and Non-VPN Traffic, page 16-4](#)

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.

TCP MSS for VPN and Non-VPN Traffic

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.

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.


Intra-Interface Communication (Routed Firewall Mode)

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. The following figure shows a network where hosts on the same interface need to communicate.



The following sample configuration shows the Cisco IOS **route-map** commands used to enable policy routing in the network shown in the figure:

```

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

```

Configure the MAC Address, MTU, and TCP MSS

Before You Begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration > Device Setup > Interface Settings > Interfaces**.

Step 2 Choose the interface row, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **Advanced** tab.

Step 4 To set the MTU or to enable jumbo frame support (supported models only), enter the value in the **MTU** field, between 300 and 9198 bytes (9000 for the ASAv).

The default is 1500 bytes.



Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.

- For models that support jumbo frames in single mode—If you enter a value for any interface that is greater than 1500, then you enable jumbo frame support automatically for all interfaces. If you set the MTU for all interfaces back to a value under 1500, then jumbo frame support is disabled.
- For models that support jumbo frames in multiple mode—If you enter a value for any interface that is greater than 1500, then be sure to enable jumbo frame support in the system configuration. See [Enable Jumbo Frame Support, page 11-8](#).



Note Enabling or disabling jumbo frame support requires you to reload the ASA.

Step 5 To manually assign a MAC address to this interface, enter a MAC address in the **Active Mac Address** field in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

Step 6 If you use failover, enter the standby MAC address in the **Standby Mac Address** field. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 7 To set the TCP MSS, choose **Configuration > Firewall > Advanced > TCP Options**. Set the following options:

- Force Maximum Segment Size for TCP—Sets the maximum TCP segment size in bytes, between 48 and any maximum number. The default value is 1380 bytes. You can disable this feature by setting the bytes to 0.
- Force Minimum Segment Size for TCP—Overrides the maximum segment size to be no less than the number of bytes you set, between 48 and any maximum number. This feature is disabled by default (set to 0).

Step 8 For **Secure Group Tagging**, see the TrustSec chapter in the firewall configuration guide.

Examples

The following example enables jumbo frames, increases the MTU on all interfaces, and disables the TCP MSS for non-VPN traffic (by setting the TCP MSS to 0, which means there is no limit):

```
jumbo frame-reservation
mtu inside 9198
mtu outside 9198
sysopt connection tcpmss 0
```


The following example enables jumbo frames, increases the MTU on all interfaces, and changes the TCP MSS for VPN traffic to 9078 (the MTU minus 120):

```
jumbo frame-reservation
mtu inside 9198
mtu outside 9198
sysopt connection tcpmss 9078
```

Allow 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.

Procedure

-
Step 1	To enable interfaces on the same security level to communicate with each other, from the Configuration > Interfaces pane, check Enable traffic between two or more interfaces which are configured with same security level .
Step 2	To enable communication between hosts connected to the same interface, check Enable traffic between two or more hosts connected to the same interface .
-

Monitoring the ARP and MAC Address Table

- **Monitoring > Interfaces > ARP Table**
Displays the ARP table, including static and dynamic entries. The ARP table includes entries that map a MAC address to an IP address for a given interface.
- **Monitoring > Interfaces > MAC Address Table**
Shows the static and dynamic MAC address entries.



Traffic Zones

You can assign multiple interfaces to a *traffic zone*, which lets traffic from an existing flow exit or enter the ASA on any interface within the zone. This capability allows Equal-Cost Multi-Path (ECMP) routing on the ASA as well as external load balancing of traffic to the ASA across multiple interfaces.

- [About Traffic Zones, page 17-1](#)
- [Prerequisites for Traffic Zones, page 17-7](#)
- [Guidelines for Traffic Zones, page 17-8](#)
- [Configure a Traffic Zone, page 17-9](#)
- [Monitoring Traffic Zones, page 17-10](#)
- [Example for Traffic Zones, page 17-12](#)
- [History for Traffic Zones, page 17-14](#)

About Traffic Zones

This section describes how you should use traffic zones in your network.

- [Non-Zoned Behavior, page 17-2](#)
- [Why Use Zones?, page 17-2](#)
- [Per-Zone Connection and Routing Tables, page 17-4](#)
- [ECMP Routing, page 17-4](#)
- [Interface-Based Security Policy, page 17-5](#)
- [Supported Services for Traffic Zones, page 17-6](#)
- [Security Levels, page 17-6](#)
- [Primary and Current Interface for the Flow, page 17-6](#)
- [Joining or Leaving a Zone, page 17-6](#)
- [Intra-Zone Traffic, page 17-6](#)
- [To- and From-the-Box Traffic, page 17-7](#)
- [Overlapping IP Addresses Within a Zone, page 17-7](#)

Non-Zoned Behavior

The Adaptive Security Algorithm takes into consideration the state of a packet when deciding to permit or deny the traffic. One of the enforced parameters for the flow is that traffic enters and exits the same interface. Any traffic for an existing flow that enters a different interface is dropped by the ASA.

Traffic zones let you group multiple interfaces together so that traffic entering or exiting *any* interface in the zone fulfills the Adaptive Security Algorithm security checks.

Related Topics

[Stateful Inspection Overview, page 1-15](#)


Why Use Zones?

You can use zones to accommodate several routing scenarios.

- [Asymmetric Routing, page 17-2](#)
- [Lost Route, page 17-3](#)
- [Load Balancing, page 17-3](#)

Asymmetric Routing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to asymmetric routing on the destination network, return traffic arrived from ISP 2 on the Outside2 interface.




Non-Zoned Problem: The ASA maintains the connection tables on a per-interface basis. When the returning traffic arrives at Outside2, it will not match the connection table and will be dropped.

Zoned Solution: The ASA maintains connection tables on a per-zone basis. If you group Outside1 and Outside2 into a zone, then when the returning traffic arrives at Outside2, it will match the per-zone connection table, and the connection will be allowed.

Lost Route

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to a lost or moved route between Outside1 and ISP 1, traffic needs to take a different route through ISP 2.




Non-Zoned Problem: The connection between the inside and outside host will be deleted; a new connection must be established using a new next-best route. For UDP, the new route will be used after a single packet drop, but for TCP, a new connection has to be reestablished.

Zoned Solution: The ASA detects the lost route and switches the flow to the new path through ISP 2. Traffic will be seamlessly forwarded without any packet drops.

Load Balancing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. A second connection was established through an equal cost route through ISP 2 on Outside2.



Non-Zoned Problem: Load-balancing across interfaces is not possible; you can only load-balance with equal cost routes on one interface.

Zoned Solution: The ASA load-balances connections across up to eight equal cost routes on all the interfaces in the zone.

Per-Zone Connection and Routing Tables

The ASA maintains a per-zone connection table so that traffic can arrive on any of the zone interfaces. The ASA also maintains a per-zone routing table for ECMP support.

ECMP Routing

The ASA supports Equal-Cost Multi-Path (ECMP) routing.

- [Non-Zoned ECMP Support, page 17-4](#)
- [Zoned ECMP Support, page 17-5](#)
- [How Connections Are Load-Balanced, page 17-5](#)
- [Falling Back to a Route in Another Zone, page 17-5](#)

Non-Zoned ECMP Support

Without zones, you can have up to 3 equal cost static or dynamic routes per interface. For example, you can configure three default routes on the outside interface that specify different gateways:

```
route outside 0 0 10.1.1.2
route outside 0 0 10.1.1.3
route outside 0 0 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

ECMP is not supported across multiple interfaces, so you cannot define a route to the same destination on a different interface. The following route is disallowed when configured with any of the routes above:

```
route outside2 0 0 10.2.1.1
```

Zoned ECMP Support

With zones, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within a zone. For example, you can configure three default routes across three interfaces in the zone:

```
route outside1 0 0 10.1.1.2
route outside2 0 0 10.2.1.2
route outside3 0 0 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The ASA load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the ASA seamlessly moves the flow to a different route.

How Connections Are Load-Balanced

The ASA load balances connections across equal cost routes using a hash made from the packet 6-tuple (source and destination IP address, source and destination port, protocol, and ingress interface). Unless the route is lost, a connection will stay on the chosen interface for its duration.

Packets within a connection are not load-balanced across routes; a connection uses a single route unless that route is lost.

The ASA does not consider the interface bandwidth or other parameters when load balancing. You should make sure all interfaces within the same zone have the same characteristics such as MTU, bandwidth, and so on.

The load-balancing algorithm is not user configurable.

Falling Back to a Route in Another Zone

When a route is lost on an interface, if there are no other routes available within the zone, then the ASA will use a route from a different interface/zone. If this backup route is used, then you may experience packet drops as with non-zoned routing support.

Interface-Based Security Policy

Zones allow traffic to and from any interface in the zone, but the security policy itself (access rules, NAT, and so on) is still applied per interface, not per zone. If you configure the same security policy for all interfaces within the zone, then you can successfully implement ECMP and load balancing for that traffic. For more information about required parallel interface configuration, see [Prerequisites for Traffic Zones, page 17-7](#).

Supported Services for Traffic Zones

The following services are supported with zones:

- Access Rules
- NAT
- Service Rules, except for QoS traffic policing.
- Routing

You can also configure to- and from-the-box services listed in [To- and From-the-Box Traffic, page 17-7](#), although full zoned support is not available.

Do not configure other services (such as VPN or Botnet Traffic Filter) for interfaces in a traffic zone; they may not function or scale as expected.

**Note**

For detailed information about how to configure the security policy, see [Prerequisites for Traffic Zones, page 17-7](#).

Security Levels

The first interface that you add to a zone determines the security level of the zone. All additional interfaces must have the same security level. To change the security level for interfaces in a zone, you must remove all but one interface, and then change the security levels, and re-add the interfaces.

Primary and Current Interface for the Flow

Each connection flow is built based on the initial ingress and egress interfaces. These interfaces are the *primary* interfaces.

If a new egress interface is used because of route changes or asymmetric routing, then the new interfaces are the *current* interfaces.

Joining or Leaving a Zone

When you assign an interface to a zone, any connections on that interface are deleted. The connections must be reestablished.

If you remove an interface from a zone, any connections that have the interface as the primary interface are deleted. The connections must be reestablished. If the interface is the current interface, the ASA moves the connections back to the primary interface. The zone route table is also refreshed.

Intra-Zone Traffic

To allow traffic to *enter* one interface and *exit* another in the same zone, enable **Configuration > Device Setup > Interface Settings > Interfaces > Enable traffic between two or more hosts connected to the same interface**, which allows traffic to enter and exit the same interface, as well as **Configuration**

> Device Setup > Interface Settings > Interfaces > Enable traffic between two or more interfaces which are configured with same security level, which allows traffic between same-security interfaces. Otherwise, a flow cannot be routed between two interfaces in the same zone.

To- and From-the-Box Traffic

- You cannot add management-only or management-access interfaces to a zone.
- For management traffic on regular interfaces in a zone, only asymmetric routing on existing flows is supported; there is no ECMP support.
- You can configure a management service on only one zone interface, but to take advantage of asymmetric routing support, you need to configure it on all interfaces. Even when the configurations are parallel on all interfaces, ECMP is not supported.
- The ASA supports the following to- and from-the-box services in a zone:
 - Telnet
 - SSH
 - HTTPS
 - SNMP
 - Syslog
 - BGP

Overlapping IP Addresses Within a Zone

For non-zoned interfaces, the ASA supports overlapping IP address networks on interfaces so long as you configure NAT properly. However, overlapping networks are not supported on interfaces in the same zone.

Prerequisites for Traffic Zones

- Configure all interface parameters including the name, IP address, and security level. Note that the security level must match for all interfaces in the zone. You should plan to group together like interfaces in terms of bandwidth and other Layer 2 properties.
- Configure the following services to match on all zone interfaces:

- Access Rules—Apply the same access rule to all zone member interfaces, or use a global access rule.

For example:

```
access-list ZONE1 extended permit tcp any host WEBSERVER1 eq 80
access-group ZONE1 in interface outside1
access-group ZONE1 in interface outside2
access-group ZONE1 in interface outside3
```

- NAT—Configure the same NAT policy on all member interfaces of the zone or use a global NAT rule (in other words, use “any” to represent the zone interfaces in the NAT rule).

Interface PAT is not supported.

For example:

```
object network WEBSERVER1
  host 10.9.9.9 255.255.255.255
  nat (inside,any) static 209.165.201.9
```



Note

When you use interface-specific NAT and PAT pools, the ASA cannot switch connections over in case of the original interface failure.

If you use interface-specific PAT pools, multiple connections from the same host might load-balance to different interfaces and use different mapped IP addresses. Internet services that use multiple concurrent connections may not work correctly in this case.

- Service Rules—Use the global service policy, or assign the same policy to each interface in a zone.

QoS traffic policing is not supported.

For example:

```
service-policy outside_policy interface outside1
service-policy outside_policy interface outside2
service-policy outside_policy interface outside3
```



Note

For VoIP inspections, zone load balancing can cause increased out-of-order packets. This situation can occur because later packets might reach the ASA before earlier packets that take a different path. Symptoms of out-of-order packets include:

- Higher memory utilization at intermediate nodes (firewall and IDS) and the receiving end nodes if queuing is used.
- Poor video or voice quality.

To mitigate these effects, we recommend that you use IP addresses only for load distribution for VoIP traffic.

- Configure routing with ECMP zone capabilities in mind.

Guidelines for Traffic Zones

Firewall Mode

Supported in routed firewall mode only. Does not support transparent firewall mode.

Failover

- You cannot add the failover or state link to a zone.
- In Active/Active failover mode, you can assign an interface in each context to an asymmetrical routing (ASR) group. This service allows traffic returning on a similar interface on the peer unit to be restored to the original unit. You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group. See [Configure Support for Asymmetrically Routed Packets \(Active/Active Mode\)](#), page 7-30 for more information about ASR groups.

- Only the primary interfaces for each connection are replicated to the standby unit; current interfaces are not replicated. If the standby unit becomes active, it will assign a new current interface if necessary.

Clustering

- You cannot add the cluster control link to a zone.

Additional Guidelines

- You can create a maximum of 256 zones.
- You can add the following types of interfaces to a zone:
 - Physical
 - VLAN
 - EtherChannel
 - Redundant
- You cannot add the following types of interfaces:
 - Management-only
 - Management-access
 - Failover or state link
 - Cluster control link
 - Member interfaces in an EtherChannel or redundant interface
 - VNI; also, if a regular data interface is marked as nve-only, it cannot be a member of a zone.
- An interface can be a member of only one zone.
- You can include up to 8 interfaces per zone.
- For ECMP, you can add up to 8 equal cost routes per zone, across all zone interfaces. You can also configure multiple routes on a single interface as part of the 8 route limit.

Configure a Traffic Zone

Configure a named zone, and assign interfaces to the zone.

Procedure

-
Step 1	Choose Configuration > Device Setup > Interface Settings > Zones , and click Add . You can alternately assign an interface to a zone from the Configuration > Device Setup > Interface Settings > Interfaces > Add Interface dialog box.
Step 2	Name the zone with a name up to 48 characters in length.
Step 3	Add one or more interfaces to the Member area. Ensure all interfaces have the same security level.
Step 4	Click Apply .
-

Monitoring Traffic Zones

This section describes how to monitor traffic zones.

- [Zone Information, page 17-10](#)
- [Zone Connections, page 17-10](#)
- [Zone Routing, page 17-11](#)

Zone Information

- **show zone [name]**

Shows zone ID, context, security level, and members.

See the following output for the **show zone** command:

```
ciscoasa# show zone outside-zone

Zone: zone-outside id: 2
Security-level: 0
Context: test-ctx
Zone Member(s) : 2
  outside1      GigabitEthernet0/0
  outside2      GigabitEthernet0/1
```

- **show nameif zone**

Shows the interface names and zone names.

See the following output for the **show nameif zone** command:

```
ciscoasa# show nameif zone

Interface      Name      zone-name  Security
GigabitEthernet0/0  inside-1  inside-zone  100
GigabitEthernet0/1.21  inside    inside-zone  100
GigabitEthernet0/1.31  4         0
GigabitEthernet0/2    outside   outside-zone  0
Management0/0      lan       0
```

Zone Connections

- **show conn [long | detail] [zone zone_name [zone zone_name] [...]]**

The **show conn zone** command displays connections for a zone. The **long** and **detail** keywords show the primary interface on which the connection was built and the current interface used to forward the traffic.

See the following output for the **show conn long zone** command:

```
ciscoasa# show conn long zone zone-inside zone zone-outside

TCP outside-zone:outside1(outside2): 10.122.122.1:1080 inside-zone:inside1(inside2):
10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

- **show asp table zone**
Shows the accelerated security path tables for debugging purposes.
- **show local-host [zone zone_name [zone zone_name] [...]]**

Shows the network states of local hosts within a zone.

See the following output for the **show local-host zone** command. The primary interface is listed first, and the current interface is in parentheses.

```
ciscoasa# show local-host zone outside-zone

Zone:outside-zone: 4 active, 5 maximum active, 0 denied
local host: <10.122.122.1>,
    TCP flow count/limit = 3/unlimited
    TCP embryonic count to host = 0
    TCP intercept watermark = unlimited
    UDP flow count/limit = 0/unlimited

Conn:
    TCP outside-zone:outside1(outside2): 10.122.122.1:1080
    inside-zone:inside1(inside2): 10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

Zone Routing

- **show route zone**

Shows the routes for zone interfaces.

See the following output for the **show route zone** command:

```
ciscoasa# show route zone

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 not set

S    192.168.105.1 255.255.255.255 [1/0] via 172.16.1.1, outside-zone:outside1
C    192.168.212.0 255.255.255.0 is directly connected, lan-zone:inside,
C    172.16.1.0 255.255.255.0 is directly connected, wan-zone:outside2
S    10.5.5.0 255.255.255.0 [1/0] via 172.16.1.1, wan-zone:outside2
O    10.2.2.1 255.255.255.255 [110/11] via 192.168.212.3, 2:09:24, lan-zone:inside
O    10.1.1.1 255.255.255.255 [110/11] via 192.168.212.2, 2:09:24, lan-zone:inside
```

- **show asp table routing**

Shows the accelerated security path tables for debugging purposes, and shows the zone associated with each route.

See the following output for the **show asp table routing** command:

```
ciscoasa# show asp table routing
route table timestamp: 60
in   255.255.255.255 255.255.255.255 identity
in   10.1.0.1       255.255.255.255 identity
in   10.2.0.1       255.255.255.255 identity
in   10.6.6.4       255.255.255.255 identity
in   10.4.4.4       255.255.255.255 via 10.4.0.10 (unresolved, timestamp: 49)
in   172.0.0.67     255.255.255.255 identity
in   172.0.0.0      255.255.255.0   wan-zone:outside2
in   10.85.43.0     255.255.255.0   via 10.4.0.3 (unresolved, timestamp: 50)
in   10.85.45.0     255.255.255.0   via 10.4.0.20 (unresolved, timestamp: 51)
```


```

in 192.168.0.0      255.255.255.0  mgmt
in 192.168.1.0      255.255.0.0    lan-zone:inside
out 255.255.255.255 255.255.255.255 mgmt
out 172.0.0.67      255.255.255.255 mgmt
out 172.0.0.0       255.255.255.0   mgmt
out 10.4.0.0        240.0.0.0       mgmt
out 255.255.255.255 255.255.255.255 lan-zone:inside
out 10.1.0.1        255.255.255.255 lan-zone:inside
out 10.2.0.0        255.255.0.0    lan-zone:inside
out 10.4.0.0        240.0.0.0    lan-zone:inside

```

Example for Traffic Zones

The following example assigns 4 VLAN interfaces to the outside zone, and configures 4 equal cost default routes. PAT is configured for the inside interface, and a web server is available on a DMZ interface using static NAT.



```

interface gigabitethernet0/0
  no shutdown
  description outside switch 1
interface gigabitethernet0/1
  no shutdown
  description outside switch 2

interface gigabitethernet0/2
  no shutdown
  description inside switch

zone outside

interface gigabitethernet0/0.101
  vlan 101
  nameif outside1

```

373595

```
security-level 0
ip address 209.165.200.225 255.255.255.224
zone-member outside
no shutdown

interface gigabitethernet0/0.102
vlan 102
nameif outside2
security-level 0
ip address 209.165.201.1 255.255.255.224
zone-member outside
no shutdown

interface gigabitethernet0/1.201
vlan 201
nameif outside3
security-level 0
ip address 198.51.100.1 255.255.255.0
zone-member outside
no shutdown

interface gigabitethernet0/1.202
vlan 202
nameif outside4
security-level 0
ip address 203.0.113.1 255.255.255.0
zone-member outside
no shutdown

interface gigabitethernet0/2.301
vlan 301
nameif inside
security-level 100
ip address 192.168.9.1 255.255.255.0
no shutdown

interface gigabitethernet0/2.302
vlan 302
nameif dmz
security-level 50
ip address 10.3.5.1 255.255.255.0
no shutdown

# Static NAT for DMZ web server on any destination interface
object network WEBSERVER
host 10.3.5.9 255.255.255.255
nat (dmz,any) static 209.165.202.129 dns

# Dynamic PAT for inside network on any destination interface
object network INSIDE
subnet 192.168.9.0 255.255.255.0
nat (inside,any) dynamic 209.165.202.130

# Global access rule for DMZ web server
access-list WEB-SERVER extended permit tcp any host WEBSERVER eq 80
access-group WEB-SERVER global

# 4 equal cost default routes for outside interfaces
route outside1 0 0 209.165.200.230
route outside2 0 0 209.165.201.10
route outside3 0 0 198.51.100.99
route outside4 0 0 203.0.113.87
# Static routes for NAT addresses - see redistribute static command
route dmz 209.165.202.129 255.255.255.255 10.3.5.99
```

```

route inside 209.165.202.130 255.255.255.255 192.168.9.99

# The global service policy
class-map inspection_default
  match default-inspection-traffic
policy-map type inspect dns preset_dns_map
  parameters
    message-length maximum client auto
    message-length maximum 512
    dns-guard
    protocol-enforcement
    nat-rewrite
policy-map global_policy
  class inspection_default
    inspect dns preset_dns_map
    inspect ftp
    inspect h323 h225 _default_h323_map
    inspect h323 ras _default_h323_map
    inspect ip-options _default_ip_options_map
    inspect netbios
    inspect rsh
    inspect rtsp
    inspect skinny
    inspect esmtp _default_esmtp_map
    inspect sqlnet
    inspect sunrpc
    inspect tftp
    inspect sip
    inspect xdmcp
service-policy global_policy global

```

History for Traffic Zones

Feature Name	Platform Releases	Description
Traffic Zones	9.3(2)	You can group interfaces together into a traffic zone to accomplish traffic load balancing (using Equal Cost Multi-Path (ECMP) routing), route redundancy, and asymmetric routing across multiple interfaces. Note You cannot apply a security policy to a named zone; the security policy is interface-based. When interfaces in a zone are configured with the same access rule, NAT, and service policy, then load-balancing and asymmetric routing operate correctly. We introduced or modified the following screens: Configuration > Device Setup > Interface Parameters > Zones Configuration > Device Setup > Interface Parameters > Interfaces



PART 4

Basic Settings



Basic Settings

This chapter describes how to configure basic settings on the ASA that are typically required for a functioning configuration.

- [Set the Hostname, Domain Name, and the Enable and Telnet Passwords, page 18-1](#)
- [Recover Enable and Telnet Passwords, page 18-2](#)
- [Set the Date and Time, page 18-7](#)
- [Configure the Master Passphrase, page 18-8](#)
- [Configure the DNS Server, page 18-10](#)
- [Adjust ASP \(Accelerated Security Path\) Performance and Behavior, page 18-12](#)
- [Monitoring the DNS Cache, page 18-13](#)
- [History for Basic Settings, page 18-14](#)

Set the Hostname, Domain Name, and the Enable and Telnet Passwords

Before you set the hostname, domain name, and the enable and Telnet passwords, check the following requirements:

Before You Begin

- In multiple context mode, you can configure the hostname and domain name in both the system and context execution spaces.
- For the enable and Telnet passwords, set them in each context; they are not available in the system. When you session to the ASASM from the switch in multiple context mode, the ASASM uses the login password you set in the admin context.
- To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

To set the hostname, domain name, and the enable and Telnet passwords, perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Device Name/Password**.
- Step 2** Enter the hostname. The default hostname is “ciscoasa.”

The hostname appears in the command line prompt, and if you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The hostname is also used in syslog messages.

For multiple context mode, the hostname that you set in the system execution space appears in the command line prompt for all contexts. The hostname that you optionally set within a context does not appear in the command line; it can be used for a banner.

Step 3 Enter the domain name. The default domain name is default.domain.invalid.

The ASA appends the domain name as a suffix to unqualified names. For example, if you set the domain name to “example.com” and specify a syslog server by the unqualified name of “jupiter,” then the ASA qualifies the name to “jupiter.example.com.”

Step 4 Change the privileged mode (enable) password. The default password is blank.

The enable password lets you enter privileged EXEC mode if you do not configure enable authentication.

The enable password also lets you log into ASDM with a blank username if you do not configure HTTP authentication.

a. Check the **Change the privileged mode password** check box.

b. Enter the old password (the default password is blank), new password, and then confirm the new password.

Step 5 Set the login password for Telnet access. There is no default password.

The login password is used for Telnet access when you do not configure Telnet authentication. You also use this password when accessing the ASASM from the switch with the **session** command.

a. Check the **Change the password to access the console of the security appliance** check box.

b. Enter the old password (for a new ASA, leave this field blank), new password, then confirm the new password.

Step 6 Click **Apply** to save your changes.

Recover Enable and Telnet Passwords

If you forget the enable or Telnet passwords, you can recover them. The procedure differs by device type. You must perform the task using the CLI.

Recover Passwords on the ASA

To recover passwords for the ASA, perform the following steps:

Procedure

Step 1 Connect to the ASA console port.

Step 2 Power off the ASA, 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. See the command reference for more information about the configuration register.

- Step 17** Save the new passwords to the startup configuration by entering the following command:

```
ciscoasa(config)# copy running-config startup-config
```

Recover Passwords on the ASA 5506-X, ASA 5508-X, and 5516-X

To recover passwords for the ASA 5506-X, ASA 5508-X, and ASA 5516-X, perform the following steps:

Procedure

- Step 1** Connect to the ASA console port.
- Step 2** Power off the ASA, 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
```

You must reset or power cycle for new config to take effect

The ASA displays the current configuration register value and a list of configuration options. Record the current configuration register value, so you can restore it later.

```
Configuration Register: 0x00000041
```

```
Configuration Summary
[ 0 ] password recovery
[ 1 ] display break prompt
[ 2 ] ignore system configuration
[ 3 ] auto-boot image in disks
[ 4 ] console baud: 9600
boot: ..... auto-boot index 1 image in disks
```

- Step 5** Reload the ASA by entering the following command:

```
rommon #2> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading disk0:/asa932-226-k8.bin... Booting...Loading...
```

The ASA loads the default configuration instead of the startup configuration.

- Step 6** Access the privileged EXEC mode by entering the following command:

```
ciscoasa# enable
```

- Step 7** When prompted for the password, press **Enter**.

The password is blank.

- Step 8** Load the startup configuration by entering the following command:

```
ciscoasa# copy startup-config running-config
```

- Step 9** Access the global configuration mode by entering the following command:

```
ciscoasa# configure terminal
```

- Step 10** 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 11** Load the default configuration by entering the following command:

```
ciscoasa(config)# no config-register
```

The default configuration register value is 0x1. See the command reference for more information about the configuration register.

- Step 12** Save the new passwords to the startup configuration by entering the following command:

```
ciscoasa(config)# copy running-config startup-config
```

Recover Passwords or Images on the ASAv

To recover passwords or images on the ASAv, perform the following steps:

Procedure

- Step 1** Copy the running configuration to a backup file on the ASAv:

```
copy running-config filename
```

Example:

```
ciscoasa# copy running-config backup.cfg
```

- Step 2** Restart the ASAv:

```
reload
```

- Step 3** From the GNU GRUB menu, press the down arrow, choose the **<filename> with no configuration load** option, then press **Enter**. The filename is the default boot image filename on the ASAv. The default boot image is never automatically booted through the **fallback** command. Then load the selected boot image.

```
GNU GRUB version 2.0(12)4
bootflash:/asa100123-20-smp-k8.bin
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

Example:

```
GNU GRUB version 2.0(12)4
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

- Step 4** Copy the backup configuration file to the running configuration.

```
copy filename running-config
```

Example:

```
ciscoasa (config)# copy backup.cfg running-config
```

- Step 5** Reset the password.

```
enable password
```

Example:

```
ciscoasa(config)# enable password cisco123
```

Step 6 Save the new configuration.

```
write memory
```

Example:

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

Disable Password Recovery



Note

You cannot disable password recovery on the ASAv.

To disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the ASA, perform the following steps.

Before You Begin

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.

Procedure

Step 1 Disable password recovery.

```
no service password-recovery
```

Example:

```
ciscoasa (config)# no service password-recovery
```

Set the Date and Time

**Note**

Do not set the date and time for the ASASM; it receives these settings from the host switch.

Set the Date and Time Using an NTP Server

NTP is used to implement a hierarchical system of servers that provide a precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure multiple NTP servers. The ASA chooses the server with the lowest stratum—a measure of how reliable the data is.

Time derived from an NTP server overrides any time set manually.

Before You Begin

In multiple context mode, you can set the time in the system configuration only.

To set the date and time using an NTP server, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Setup > System Time > NTP**.
- Step 2** Click **Add** to display the **Add NTP Server Configuration** dialog box.
- Step 3** Enter the NTP server IP address.
- Step 4** Check the **Preferred** check box to set this server as a preferred server. NTP uses an algorithm to determine which server is the most accurate and synchronizes to it. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the ASA uses the more accurate one.
- Step 5** Choose the interface from the drop-down list. This setting specifies the outgoing interface for NTP packets. If the interface is blank, then the ASA uses the default admin context interface according to the routing table. Choose **None** (the default interface) for stability to change the admin context (and the available interfaces).
- Step 6** Choose the key number from the drop-down list. This setting specifies the key ID for this authentication key, which enables you to use MD5 authentication to communicate with the NTP server. The NTP server packets must also use this key ID. If you have previously configured a key ID for another server, you can select it from the list; otherwise, enter a number between 1 and 4294967295.
- Step 7** Check the **Trusted** check box to set this authentication key as a trusted key, which is required for authentication to succeed.
- Step 8** Enter the key value to set the authentication key, which is a string that can be up to 32 characters long.
- Step 9** Reenter the key value to make sure that you enter it correctly twice.
- Step 10** Click **OK**.
- Step 11** Check the **Enable NTP authentication** check box to turn on NTP authentication.
- Step 12** Click **Apply** to save your changes.

Set the Date and Time Manually

Before setting the date and time manually, check the following requirements:

Before You Begin

In multiple context mode, you can set the time in the system configuration only.

To set the date and time manually, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Setup > System Time > Clock**.
- Step 2** Choose the time zone from the drop-down list. This setting specifies the time zone as GMT plus or minus the appropriate number of hours. If you select the Eastern Time, Central Time, Mountain Time, or Pacific Time zone, then the time adjusts automatically for daylight savings time, from 2:00 a.m. on the second Sunday in March to 2:00 a.m. on the first Sunday in November.
- 

Note Changing the time zone on the ASA may drop the connection to intelligent SSMs.
-
- Step 3** Click the **Date** drop-down list to display a calendar. Then find the correct date using the following methods:
- Click the name of the month to display a list of months, then click the desired month. The calendar updates to that month.
 - Click the year to change the year. Use the up and down arrows to scroll through the years, or enter a year in the entry field.
 - Click the arrows to the right and left of the month and year to scroll the calendar forward and backward one month at a time.
 - Click a day on the calendar to set the date.
- Step 4** Enter the time manually in hours, minutes, and seconds.
- Step 5** Click **Update Display Time** to update the time shown in the bottom right corner of the ASDM pane. The current time updates automatically every ten seconds.
-

Configure 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

**Note**

If failover is enabled but no failover shared key is set, an error message appears if you change the master passphrase, informing you that you must enter a failover shared key to protect the master passphrase changes from being sent as plain text.

Choose **Configuration > Device Management > High Availability > Failover**, enter any character in the Shared Key field or 32 hexadecimal numbers (0-9A-Fa-f) if a failover hexadecimal key is selected, except a backspace. Then click **Apply**.

Add or Change the Master Passphrase

Before you add or change the master passphrase, check the following requirements:

Before You Begin

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:

Procedure

-
- Step 1** Choose one of the following options:
- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.
 - In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.
- Step 2** Check the **Advanced Encryption Standard (AES) password encryption** check box.
- If no master passphrase is in effect, a warning message appears when you click **Apply**. You can click **OK** or **Cancel** to continue.
- If you later disable password encryption, all existing encrypted passwords are left unchanged, and as long as the master passphrase exists, the encrypted passwords will be decrypted, as required by the application.
- Step 3** Check the **Change the encryption master passphrase** check box to enable you to enter and confirm your new master passphrases. By default, they are disabled.
- Your new master passphrase must be between 8 and 128 characters long.
- If you are changing an existing passphrase, you must enter the old passphrase before you can enter a new one.
- Leave the **New** and **Confirm master passphrase** fields blank to delete the master passphrase.
- Step 4** Click **Apply**.
-

Disable 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.

Before You Begin

- You must know the current master passphrase to disable it.
- 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:

Procedure

-
- Step 1** Choose one of the following options:
- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.
 - In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.
- Step 2** Check the **Advanced Encryption Standard (AES) password encryption** check box.
- If no master passphrase is in effect, a warning statement appears when you click **Apply**. Click **OK** or **Cancel** to continue.
- Step 3** Check the **Change the encryption master passphrase** check box.
- Step 4** Enter the old master passphrase in the **Old master passphrase** field. You must provide the old master passphrase to disable it.
- Step 5** Leave the **New master passphrase** and the **Confirm master passphrase** fields empty.
- Step 6** Click **Apply**.
-

Configure the DNS Server

You need to configure DNS servers so that the ASA can resolve host names to IP addresses. You also must configure DNS servers to use fully qualified domain names (FQDN) network objects in access rules.

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.

Before You Begin

Make sure that you configure the appropriate routing and access rules for any interface on which you enable DNS domain lookup so you can reach the DNS server.

To configure the DNS server, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Management > DNS > DNS Client**.
- Step 2** Ensure that at DNS lookup is enabled on at least one interface. In the **DNS Lookup** interface list, below the DNS server group table, click in the DNS Enabled column and select **True** to enable lookup on the interface.
- Step 3** Choose one of the following options in the **DNS Setup** area:
- **Configure one DNS server group.**
 - **Configure multiple DNS server groups.**
- Step 4** Do one of the following:
- Select a DNS group and click **Edit**.
 - If you elected to configure multiple DNS groups, click **Add** to add a new group. Enter a group name.
- Step 5** Configure the DNS server group.
- a. Enter the IP addresses of the configured servers, then click **Add**.
You can add up to six DNS servers. The ASA tries each DNS server in order until it receives a response. Use the **Move Up/Move Down** buttons to put the servers in priority order.
 - b. Enter the number of seconds (between 1 and 30) to wait before trying the next DNS server in the list in the **Other Settings** area. The default is 2 seconds. Each time the ASA retries the list of servers, the timeout time doubles.
 - c. Enter a DNS domain name for the group of configured servers.
 - d. Click **OK**.
- Step 6** If you have multiple groups, select the one you want to use and click **Set Active**. This server group will be used for DNS requests.
- Step 7** Check the **Enable DNS Guard on all interfaces** check box to enforce one DNS response per query.
You can also set DNS Guard when configuring DNS inspection. For a given interface, the DNS Guard setting configured in DNS inspection takes precedence over this global setting. By default, DNS inspection is enabled on all interfaces with DNS Guard enabled.
- Step 8** Click **Apply** to save your changes.
-

Adjust ASP (Accelerated Security Path) Performance and Behavior

The ASP is an implementation layer that puts your policies and configurations into action. It is not of direct interest except during troubleshooting with the Cisco Technical Assistance Center. However, there are a few behaviors related to performance and reliability that you can adjust.

- [Choose a Rule Engine Transactional Commit Model, page 18-12](#)
- [Enable ASP Load Balancing, page 18-13](#)

Choose a Rule Engine Transactional Commit Model

By default, when you change a rule-based policy (such as access rules), the changes become effective immediately. However, this immediacy comes with a slight cost in performance. The performance cost is more noticeable for very large rule lists in a high connections-per-second environment, for example, when you change a policy with 25,000 rules while the ASA is handling 18,000 connections per second.

The performance is affected because the rule engine compiles rules to enable faster rule lookup. By default, the system also searches uncompiled rules when evaluating a connection attempt so that new rules can be applied; because the rules are not compiled, the search takes longer.

You can change this behavior so that the rule engine uses a transactional model when implementing rule changes, continuing to use the old rules until the new rules are compiled and ready for use. With the transactional model, performance should not drop during the rule compilation. The following table clarifies the behavioral difference.

Model	Before Compilation	During Compilation	After Compilation
Default	Matches old rules.	Match new rules. (The rate for connections per second decreases.)	Matches new rules.
Transactional	Matches old rules.	Match old rules. (The rate for connections per second is unaffected.)	Matches new rules.

An additional benefit of the transactional model is that, when replacing an ACL on an interface, there is no gap between deleting the old ACL and applying the new one. This feature reduces the chances that acceptable connections may be dropped during the operation.

**Tip**

If you enable the transactional model for a rule type, syslogs to mark the beginning and the end of the compilation are generated. These syslogs are numbered 780001 through 780004.

To enable the transactional commit model for the rule engine, choose **Configuration > Device Management > Advanced > Rule Engine** and select the desired options:

- **Access group**—Access rules applied globally or to interfaces.
- **NAT**—Network address translation rules.

Enable ASP Load Balancing

The ASP load balancing mechanism helps avoid the following issues:

- Overruns caused by sporadic traffic spikes on flows
- Overruns caused by bulk flows oversubscribing specific interface receive rings
- Overruns caused by relatively heavily overloaded interface receive rings, in which a single core cannot sustain the load.

The **asp load-balance per-packet** command allows multiple cores to work simultaneously on packets that were received from a single interface receive ring. If the system drops packets, and the **show cpu** command output is far less than 100%, then this command may help your throughput if the packets belong to many unrelated connections. The **auto** option enables the ASA to automatically switch per-packet load balancing on and off.

On ASA models with multiple cores, if you see many packet drops but CPU usage is significantly below 100%, then you should enable the load balancing option.

Choose **Configuration > Device Management > Advanced > ASP Load Balancing**, then check the **Enable ASP load balancing** check box.

Check the **Dynamically enable or disable ASP load balancing based on traffic monitoring** check box to enable ASP load balancing automatically.

Monitoring the 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.

See the following command for monitoring the DNS cache:

- **show dns-hosts**

This command shows 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.

History for Basic Settings

Feature Name	Platform Releases	Description
Master Passphrase	8.3(1)	We introduced this feature. The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. We introduced the following screens: Configuration > Device Management > Advanced > Master Passphrase. Configuration > Device Management > Device Administration > Master Passphrase.
Password Encryption Visibility	8.4(1)	We modified the show password encryption command.
Removal of the default Telnet password	9.0(2)/9.1(2)	To improve security for management access to the ASA, the default login password for Telnet was removed; you must manually set the password before you can log in using Telnet. Note The login password is only used for Telnet if you do not configure Telnet user authentication. Previously, when you cleared the password, the ASA restored the default of “cisco.” Now when you clear the password, the password is removed. The login password is also used for Telnet sessions from the switch to the ASASM (see the session command). For initial ASASM access, you must use the service-module session command, until you set a login password. We did not modify any ASDM screens.
ASP Load Balancing	9.3(2)	We introduced this feature. The ASP load balancing mechanism reduces packet drop and improves throughput by allowing multiple cores of the CPU to receive packets from an interface receive ring and work on them independently. We introduced the following screen: Configuration > Device Management > Advanced > ASP Load Balancing.



DHCP and DDNS Services

This chapter describes how to configure the DHCP server or DHCP relay as well as dynamic DNS (DDNS) update methods.

- [About DHCP and DDNS, page 19-1](#)
- [Guidelines for DHCP and DDNS Services, page 19-3](#)
- [Configure the DHCP Server, page 19-5](#)
- [Configure DDNS, page 19-9](#)
- [Monitoring DHCP and DDNS Services, page 19-11](#)
- [History for DHCP and DDNS Services, page 19-12](#)

About DHCP and DDNS

This section describes how DHCP clients and servers work together using the DHCP relay agent, as well as how DDNS integrates with DHCP.

About the DHCP Server

DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The Cisco 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.

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.

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.

The DDNS name and address mapping is held on the DHCP server in two resource records (RRs): 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.

Related Topics

- [Configure the DHCP Server, page 19-5](#)

DDNS Update Configurations

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. The DHCP server must know the fully qualified domain name (FQDN) of the client to update the PTR RR. The client provides an FQDN to the server using a DHCP option called Client FQDN.

UDP Packet Size

DDNS 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 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** and **message-length maximum server 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 allows the automatically configured length up to the configured client or server maximum. For all other DNS traffic, the message-length maximum is used.

Guidelines for DHCP and DDNS Services

This section includes guidelines and limitations that you should check before configuring DHCP and DDNS services.

Firewall Mode

Not supported in transparent firewall mode.

IPv6

Does not support IPv6 for interface-specific DHCP relay servers.

DHCP Server

- 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.
- 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

- 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.

- When the ASA relays DHCP to the DHCP server, it sources the packet with the address of the interface facing the DHCP server instead of the address facing the DHCP client (GIADDR). This address becomes problematic when combined with EasyVPN deployments because it may not be unique, and the DHCP traffic must be routed through a VPN tunnel. The ASA EasyVPN server does not support multiple peers having the same addresses. To correct this issue, the ASA should source the packet using the address to which the DHCP server will send its response (GIADDR). This address must already be unique when deploying DHCP relay.

Configure the DHCP Server

This section describes how to configure a DHCP server provided by the ASA.

-
Step 1	Enable the DHCP Server. See Enable the DHCP Server, page 19-5 .
Step 2	Configure advanced DHCP options. See Configure Advanced DHCP Options, page 19-7 .
Step 3	Configure either a DHCPv4 relay agent or a DHCPv6 relay agent. See Configure the DHCPv4 Relay Agent, page 19-8 or Configure the DHCPv6 Relay Agent, page 19-8 .
-

Enable the DHCP Server

To enable the DHCP server on an ASA interface, perform the following steps:

Procedure

-
Step 1	Choose Configuration > Device Management > DHCP > DHCP Server .
Step 2	Choose an interface, then click Edit . a. Check the Enable DHCP Server check box to enable the DHCP server on the selected interface. b. Enter the range of IP addresses from lowest to highest that is used by the DHCP server in the DHCP Address Pool field. The range of IP addresses must be on the same subnet as the selected interface and cannot include the IP address of the interface itself. c. Set the following in the Optional Parameters area: – The DNS servers (1 and 2) configured for the interface. – The WINS servers (primary and secondary) configured for the interface. – The domain name of the interface. – The time in milliseconds that the ASA will wait for an ICMP ping response on the interface. – The duration of time that the DHCP server configured on the interface allows DHCP clients to use an assigned IP address. – The interface on a DHCP client that provides DNS, WINS, and domain name information for automatic configuration if the ASA is acting as a DHCP client on a specified interface (usually outside). – Click Advanced to display the Advanced DHCP Options dialog box to configure more DHCP options. See Configure Advanced DHCP Options, page 19-7 for more information.

- d. Check the **Update DNS Clients** check box in the **Dynamic Settings for DHCP Server** area to specify that in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:
 - Check the **Update Both Records** check box to specify that the DHCP server should update both the A and PTR RRs.
 - Check the **Override Client Settings** check box to specify that DHCP server actions should override any update actions requested by the DHCP client.
 - e. Click **OK** to close the **Edit DHCP Server** dialog box.
- Step 3** Check the **Enable Auto-configuration from interface** check box in the **Global DHCP Options** area below the DHCP Server table to enable DHCP auto configuration only if the ASA is acting as a DHCP client on a specified interface (usually outside).

DHCP auto configuration enables the DHCP Server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client that is running on the specified interface. If information obtained through auto configuration is also specified manually in the **Global DHCP Options** area, the manually specified information takes precedence over the discovered information.
- Step 4** Choose the interface from the drop-down list.
- Step 5** Check the **Allow VPN override** check box to override the interface DHCP or PPPoE client WINS parameter with the VPN client parameter.
- Step 6** Enter the IP address of the primary DNS server for a DHCP client in the **DNS Server 1** field.
- Step 7** Enter the IP address of the alternate DNS server for a DHCP client in the **DNS Server 2** field.
- Step 8** Enter the DNS domain name for DHCP clients (for example, example.com) in the **Domain Name** field.
- Step 9** Enter the amount of time, in seconds, in the **Lease Length** field that the client may use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).
- Step 10** Enter the IP address of the primary WINS server for a DHCP client in the **Primary WINS Server** field.
- Step 11** Enter the IP address of the alternate WINS server for a DHCP client in the **Secondary WINS Server** field.
- Step 12** To avoid address conflicts, the ASA sends two ICMP ping packets to an address before assigning that address to a DHCP client. Enter the amount of time, in milliseconds, in the **Ping Timeout** field that the ASA waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.
- Step 13** Click **Advanced** to display the **Configuring Advanced DHCP Options** dialog box to specify additional DHCP options and their parameters. For more information, see [Configure Advanced DHCP Options, page 19-7](#).
- Step 14** You configure the DDNS update settings for the DHCP server in the **Dynamic DNS Settings for DHCP Server** area. Check the **Update DNS Clients** check box to specify that, in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:
 - Check the **Update Both Records** check box to specify that the DHCP server should update both the A and PTR RRs.
 - Check the **Override Client Settings** check box to specify that the DHCP server actions should override any update actions requested by the DHCP client.

Step 15 Click **Apply** to save your changes.

Configure Advanced DHCP Options

The ASA supports the DHCP options listed in RFC 2132, RFC 2562, and RFC 5510 to send information. You can use advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients. You can also use the DHCP automatic configuration setting to obtain these values or define them manually. When you use more than one method to define this information, it is passed to DHCP clients in the following sequence:

1. Manually configured settings.
2. Advanced DHCP options settings.
3. DHCP automatic configuration settings.

For example, you can manually define the domain name that you want the DHCP clients to receive and then enable DHCP automatic configuration. Although DHCP automatic configuration discovers the domain together with the DNS and WINS servers, the manually defined domain name is passed to DHCP clients with the discovered DNS and WINS server names, because the domain name discovered by the DHCP automatic configuration process is superseded by the manually defined domain name.

Procedure

- Step 1** Choose **Configuration > Device Management > DHCP > DHCP Server**, then click **Advanced**.
- Step 2** Choose the option code from the drop-down list. All DHCP options (1 through 255) are supported except for 1, 12, 50–54, 58–59, 61, 67, and 82.
- Step 3** Choose the options that you want to configure. Some options are standard. For standard options, the option name is shown in parentheses after the option number and the option parameters are limited to those supported by the option. For all other options, only the option number is shown and you must choose the appropriate parameters to supply with the option. For example, if you choose DHCP Option 2 (Time Offset), you can only enter a hexadecimal value for the option. For all other DHCP options, all of the option value types are available and you must choose the appropriate one.
- Step 4** Specify the type of information that the option returns to the DHCP client in the **Option Data** area. For standard DHCP options, only the supported option value type is available. For all other DHCP options, all of the option value types are available. Click **Add** to add the option to the DHCP option list. Click **Delete** to remove the option from the DHCP option list.
 - Click **IP Address** to indicate that an IP address is returned to the DHCP client. You can specify up to two IP addresses. IP Address 1 and IP Address 2 indicate an IP address in dotted-decimal notation.



Note The name of the associated IP address fields can change based on the DHCP option that you chose. For example, if you choose DHCP Option 3 (Router), the fields names change to Router 1 and Router 2.

 - Click **ASCII** to specify that an ASCII value is returned to the DHCP client. Enter an ASCII character string in the **Data** field. The string cannot include spaces.



Note

The name of the associated Data field can change based on the DHCP option that you chose. For example, if you choose DHCP Option 14 (Merit Dump File), the associated Data field names change to File Name.

- Click **Hex** to specify that a hexadecimal value is returned to the DHCP client. Enter a hexadecimal string with an even number of digits and no spaces in the **Data** field. You do not need to use a 0x prefix.



Note

The name of the associated Data field can change based on the DHCP option you chose. For example, if you choose DHCP Option 2 (Time Offset), the associated Data field becomes the Offset field.

Step 5 Click **OK** to close the **Advanced DHCP Options** dialog box.

Step 6 Click **Apply** to save your changes.

Configure 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 may 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.

Configure the DHCPv6 Relay Agent

When a DHCPv6 request enters an interface, then the ASA relays the request to all DHCPv6 global servers.

Procedure

Step 1 Choose **Configuration > Device Management > DHCP > DHCP Relay**.

Step 2 Check the check boxes for the services you want for each interface in the **DHCP Relay Agent** area.

- IPv4 > DHCP Relay Enabled.**
- IPv4 > Set Route**—Changes the default gateway address in the DHCP message from the server to that of the ASA interface that is closest to the DHCP client, which relayed the original DHCP request. This action allows the client to set its default route to point to the ASA even if the DHCP server specifies a different router. If there is no default router option in the packet, the ASA adds one containing the interface address.
- IPv6 > DHCP Relay Enabled.**

- **Trusted Interface**—Specifies a DHCP client interface that you want to trust. You can configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. You can alternatively trust all interfaces by checking the **Set dhcp relay information as trusted on all interfaces** check box (see [Step 7](#)).

- Step 3** Add one or more DHCP servers to which DHCP requests are relayed in the **Global DHCP Relay Servers** area,
- Click **Add**. The **Add Global DHCP Relay Server** dialog box appears.
 - Enter the IPv4 or IPv6 address of the DHCP server in the **DHCP Server** field.
 - Choose the interface to which the specified DHCP server is attached from the **Interface** drop-down list.
 - Click **OK**.
- The newly added global DHCP relay server appears in the **Global DHCP Relay Servers** list.
- Step 4** (Optional) In the **IPv4 Timeout** field, enter the amount of time, in seconds, allowed for DHCP address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.
- Step 5** (Optional) In the **IPv6 Timeout** field, enter the amount of time, in seconds, allowed for DHCP address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.
- Step 6** In the **DHCP Relay Interface Servers** area, add one or more interface-specific DHCP servers to which DHCP requests on a given interface are relayed:
- Click **Add**. The **Add DHCP Relay Server** dialog box appears.
 - From the **Interface** drop-down list, choose the interface connected to the DHCP clients. Note that you do not specify the egress interface for the requests, as for a Global DHCP Server; instead, the ASA uses the routing table to determine the egress interface.
 - In the **Server to...** field, enter the IPv4 address of the DHCP server, and click **Add>>**. The server is added to the right-hand list. Add up to 4 servers, if available out of the overall maximum. IPv6 is not supported for interface-specific servers.
 - Click **OK**.
- The newly added interface DHCP relay server(s) appear in the **DHCP Relay Interface Servers** list.
- Step 7** To configure all interfaces as trusted interfaces, check the **Set dhcp relay information as trusted on all interfaces** check box. You can alternatively trust individual interfaces (see [Step 2](#)).
- Step 8** Click **Apply** to save your settings.

Configure DDNS

This section describes how to configure DDNS.

To configure dynamic DNS and update the DNS server, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Management > DNS > Dynamic DNS**.
- Step 2** Click **Add** to display the **Add Dynamic DNS Update Method** dialog box.
- Step 3** Enter the name for the DDNS update method.
- Step 4** Specify the update interval between DNS update attempts configured for the update method in days, hours, minutes, and seconds.
- Choose the number of days between update attempts from 0 to 364.
 - Choose the number of hours (in whole numbers) between update attempts from 0 to 23.
 - Choose the number of minutes (in whole numbers) between update attempts from 0 to 59.
 - Choose the number of seconds (in whole numbers) between update attempts from 0 to 59.
- These units are additive. That is, if you enter 0 days, 0 hours, 5 minutes and 15 seconds, the update method tries an update every 5 minutes and 15 seconds for as long as the method is active.
- Step 5** Choose one of the following options to store server resource record updates that the DNS client updates:
- Both the A resource record and the PTR resource record.
 - The A resource records only.
- Step 6** Click **OK** to close the **Add Dynamic DNS Update Method** dialog box.
- The new dynamic DNS client settings appear.



Note When you edit an existing method, the Name field is *display-only* and shows the name of the selected method for editing.

- Step 7** Click **Add** to display the **Add Dynamic DNS Interface Settings** dialog box to add DDNS settings for each interface configured.
- Step 8** Choose the interface from the drop-down list.
- Step 9** Choose the update method assigned to the interface from the drop-down list.
- Step 10** Enter the hostname of the DDNS client.
- Step 11** Choose one of the following options to store resource record updates:
- Default (PTR Records) to specify that the client request PTR record updating by the server.
 - Both (PTR Records and A Records) to specify that the client request both the A and PTR DNS resource records by the server.
 - None to specify that the client request no updates by the server.



Note DHCP must be enabled on the selected interface for this action to take effect.

- Step 12** Click **OK** to close the **Add Dynamic DNS Interface Settings** dialog box.
- The new dynamic DNS interface settings appear.
- Step 13** Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.
-

Monitoring DHCP and DDNS Services

This section includes the procedures to monitor both DHCP and DDNS services.

Monitoring DHCP Services

See the following screens for monitoring DHCP services:

- **Monitoring > Interfaces > DHCP > DHCP Client Lease Information.**
This pane displays configured DHCP client IP addresses.
- **Monitoring > Interfaces > DHCP > DHCP Server Table**
This pane displays configured dynamic DHCP client IP addresses.
- **Monitoring > Interfaces > DHCP > DHCP Statistics**
This pane displays DHCPv4 message types, counters, values, directions, messages received, and messages sent.
- **Monitoring > Interfaces > DHCP > IPV6 DHCP Statistics**
This pane displays DHCPv6 message types, counters, values, directions, messages received, and messages sent.
- **Monitoring > Interfaces > DHCP > IPV6 DHCP Binding**
This pane displays DHCPv6 bindings.
- **Tools > Command Line Interface**
This pane sends non-interactive commands to the ASA and lists results.

Monitoring DDNS Status

See the following screens for monitoring DDNS status:

- **Tools > Command Line Interface**
This pane sends non-interactive commands to the ASA and lists results.

History for DHCP and DDNS Services

Table 19-1 History for DHCP and DDNS Services

Feature Name	Platform Releases	Description
DHCP	7.0(1)	The ASA can provide a DHCP server or DHCP relay services to DHCP clients attached to ASA interfaces. We introduced the following screens: Configuration > Device Management > DHCP > DHCP Relay. Configuration > Device Management > DHCP > DHCP Server.
DDNS	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management > DNS > DNS Client. Configuration > Device Management > DNS > Dynamic DNS.
DHCP for IPv6 (DHCPv6)	9.0(1)	Support for IPv6 was added. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP relay servers per interface (IPv4 only)	9.1(2)	You can now configure DHCP relay servers per-interface, so requests that enter a given interface are relayed only to servers specified for that interface. IPv6 is not supported for per-interface DHCP relay. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP trusted interfaces	9.1(2)	You can now configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP rebind function	9.1(4)	During the DHCP rebind phase, the client now tries to rebind to other DHCP servers in the tunnel group list. Before this release, the client did not rebind to an alternate server when the DHCP lease fails to renew. We did not modify any ASDM screens.
DHCP monitoring for IPv6	9.4(1)	You can now monitor DHCP for IPv6. We introduced the following screens: Monitoring > Interfaces > DHCP > IPV6 DHCP Statistics Monitoring > Interfaces > DHCP > IPV6 DHCP Binding.



Digital Certificates

This chapter describes how to configure digital certificates.

- [About Digital Certificates, page 20-1](#)
- [Guidelines for Digital Certificates, page 20-9](#)
- [History for Certificate Management, page 20-28](#)

About Digital Certificates

Digital certificates provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security.

For authentication using digital certificates, at least one identity certificate and its issuing CA certificate must exist on an ASA. This configuration allows multiple identities, roots, and certificate hierarchies. The ASA evaluates third-party certificates against CRLs, also called authority revocation lists, all the way from the identity certificate up the chain of subordinate certificate authorities.

Descriptions of several different types of available digital certificates follow:

- A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. A certificate that is issued by another CA certificate is called a subordinate certificate.
- CAs also issue identity certificates, which are certificates for specific systems or hosts.
- Code-signer certificates are special certificates that are used to create digital signatures to sign code, with the signed code itself revealing the certificate origin.

The local CA integrates an independent certificate authority feature on the ASA, deploys certificates, and provides secure revocation checking of issued certificates. The local CA provides a secure, configurable, in-house authority for certificate authentication with user enrollment through a website login page.



Note

CA certificates and identity certificates apply to both site-to-site VPN connections and remote access VPN connections. Procedures in this document refer to remote access VPN use in the ASDM GUI.

Digital certificates provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security.

For authentication using digital certificates, at least one identity certificate and its issuing CA certificate must exist on an ASA. This configuration allows multiple identities, roots, and certificate hierarchies. Descriptions of several different types of available digital certificates follow:

- A CA certificate is used to sign other certificates. It is self-signed and called a root certificate.
- A certificate that is issued by another CA certificate is called a subordinate certificate.

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>.

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 the Cisco 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 Cisco 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 ocsb** command. You can also make the OCSP check optional by using the **revocation-check ocsb 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 **ocsb 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 **ocsb-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 ocsb** 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, syslog messages 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.

The following figure shows that 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 20-1 The Local CA



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.

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

Certificates

If user digital certificates are configured, the ASA first validates the certificate. It does not, however, use any of the DN fields 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`.

Guidelines for Digital Certificates

This section includes guidelines and limitations that you should check before configuring digital certificates.

Context Mode Guidelines

- Supported in single context mode only for third-party CAs.

Failover Guidelines

- Does not support replicating sessions in Stateful Failover.
- Does not support failover for local CAs.

IPv6 Guidelines

Does not support IPv6.

Local CA Certificates

- 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.
- 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.
- 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.

The following example shows a base 64 encoded local CA certificate:

```
MIIXlwIBAZCCF1EGCSqGSIb3DQEHAaCCF0IEghc+MIIXOjCCFzYGCSqGSIb3DQEHBqCCFycwghcAgEAMIIXHAYJKoZIhvcNAQcBMBsGCiqGSIb3DQEMAQMwDQQIjph4SxJoyTgCAQGAghbw3v4bFy+GGG2dJnB4OLphsUM+IG3S
DOiDwZG9n1SvtMieoxd7Hxknxbum06JDrujWktHBIqkrm+td34qlNE1iGeP2YC94/NQ2z+4kS+uZzwcRh1lKEZ
TS1E4L0fSaC3uMTxJq2NUHYWmoc8pi4CIeLj3h7VVMY6qbx2AC8I+q57+QG5vG515Hi5imwtYfaWwPEdPQxaWZ
```

```
PrzoG1J8BFqdPa1jBGhAzzuSmElm3j/2dQ3Atro1G9nIsRHgV39fcBgwz4fEabHG7/Vanb+fj81d5n1OiJjDYY  
bP86tvbZ2yOVZR6aKFVI0b2AfCr6PbwfC9U8Z/aF3BCyM2sN2xPJrXva94CaYrqyotZdAkSYA5KWSyEcgdqmu  
BeGDKOncTknfgy0XM+fG5rb3qAXy1GkjyFI5Bm9Do6RUR0oG1DSrQrKeq/hj...
```

END OF CERTIFICATE

SCEP Proxy Support

- 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.

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.

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°3é¼p0 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
```

Add or Install a CA Certificate

To add or install a CA certificate, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Remote Access VPN > Certificate Management > CA Certificates**.
- Step 2** Click **Add**.
The **Install Certificate** dialog box appears.
- Step 3** Click the **Install from a file** radio button to add a certificate configuration from an existing file (this is the default setting).
- Step 4** Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.
- Step 5** The **Certificate Installation** dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.
- Step 6** Click the **Paste certificate in PEM format** radio button to enroll manually.
- Step 7** Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided, then click **Install Certificate**.
- Step 8** The **Certificate Installation** dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.
- Step 9** Click the **Use SCEP** radio button to enroll automatically. The ASA contacts the CA using SCEP, obtains the certificates, and installs them on the device. To use SCEP, you must enroll with a CA that supports SCEP, and you must enroll via the Internet. Automatic enrollment using SCEP requires that you provide the following information:
 - The path and file name of the certificate to be automatically installed.
 - The maximum number of minutes to retry certificate installation. The default is one minute.
 - The number of retries for installing a certificate. The default is zero, which indicates unlimited retries within the retry period.

Step 10 Click **More Options** to display additional configuration options for new and existing certificates. The **Configuration Options for CA Certificates** pane appears.

Step 11 To change an existing CA certificate configuration, select it, then click **Edit**.

Step 12 To remove a CA certificate configuration, select it, then click **Delete**.

**Note**

After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Step 13 Click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:

- The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trust points. The values apply to both available and pending status.
- The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Configure CA Certificates for Revocation

To configure CA certificates for revocation, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the **Install Certificates** dialog box. Then click **More Options**.
- Step 2** Click the **Revocation Check** tab.
- Step 3** Click the **Do not check certificates for revocation** radio button to disable revocation checking of certificates.
- Step 4** Click the **Check certificates for revocation** radio button to select one or more revocation checking methods (CRL or OCSP).
- Step 5** Click **Add** to move a revocation method to the right and make it available. Click **Move Up** or **Move Down** to change the method order.
- The methods you choose are implemented in the order in which you add them. If a method returns an error, the next revocation checking method activates.
- Step 6** Check the **Consider certificate valid if revocation checking returns errors** check box to ignore revocation checking errors during certificate validation.
- Step 7** Click **OK** to close the **Revocation Check** tab.

Configure CRL Retrieval Policy

To configure the CRL retrieval policy, perform the following steps:

Procedure

-
Step 1	Choose Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add to display the Install Certificates dialog box. Then click More Options .
Step 2	Check the Use CRL Distribution Point from the certificate check box to direct revocation checking to the CRL distribution point from the certificate being checked.
Step 3	Check the Use Static URLs configured below check box to list specific URLs to be used for CRL retrieval. The URLs you select are implemented in the order in which you add them. If an error occurs with the specified URL, the next URL in order is taken.
Step 4	Click Add in the Static Configuration area. The Add Static URL dialog box appears.
Step 5	Enter the static URL to use for distributing the CRLs, then click OK . The URL that you entered appears in the Static URLs list.
Step 6	Click OK to close this dialog box.
-

Configure CRL Retrieval Methods

To configure CRL retrieval methods, perform the following steps:

Procedure

-
Step 1	Choose Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add to display the Install Certificates dialog box. Then click More Options .
Step 2	Click the CRL Retrieval Methods tab in the Configuration Options for CA Certificates pane.
Step 3	Choose one of the following three retrieval methods: • To enable LDAP for CRL retrieval, check the Enable Lightweight Directory Access Protocol (LDAP) check box. With LDAP, CRL retrieval starts an LDAP session by connecting to a named LDAP server, accessed by a password. The connection is on TCP port 389 by default. Enter the following required parameters: – Name – Password – Confirm Password – Default Server (server name) – Default Port (389) • To enable HTTP for CRL retrieval, check the Enable HTTP check box.
Step 4	Click OK to close this tab.
-

Configure OCSP Rules

This section describes how to configure OCSP rules.

Before You Begin

Make sure that you have configured a certificate map before you try to add OCSP rules. If a certificate map has not been configured, an error message appears.

To configure OCSP rules for obtaining revocation status of an X.509 digital certificate, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the **Install Certificates** dialog box. Then click **More Options**.
 - Step 2** Click the **OCSP Rules** tab in the **Configuration Options for CA Certificates** pane.
 - Step 3** Choose the certificate map to match to this OCSP rule. Certificate maps match user permissions to specific fields in a certificate. The name of the CA that the ASA uses to validate responder certificates appears in the **Certificate** field. The priority number for the rule appears in the **Index** field. The URL of the OCSP server for this certificate appears in the **URL** field.
 - Step 4** Click **Add**.
The **Add OCSP Rule** dialog box appears.
 - Step 5** Choose the certificate map to use from the drop-down list.
 - Step 6** Choose the certificate to use from the drop-down list.
 - Step 7** Enter the priority number for the rule.
 - Step 8** Enter the URL of the OCSP server for this certificate.
 - Step 9** When you are done, click **OK** to close this dialog box.
The newly added OCSP rule appears in the list.
 - Step 10** Click **OK** to close this tab.
-

Configure Advanced CRL and OCSP Settings

To configure additional CRL and OCSP settings, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Site-to-Site VPN > Certificate Management > CA Certificates > Add** to display the **Install Certificates** dialog box. Then click **More Options**.
 - Step 2** Click the **Advanced** tab in the **Configuration Options for CA Certificates** pane.
 - Step 3** Enter the number of minutes between cache refreshes in the **CRL Options** area. The default is 60 minutes. The range is 1-1440 minutes. To avoid having to retrieve the same CRL from a CA repeatedly, the ASA can store retrieved CRLs locally, which is called CRL caching. The CRL cache capacity varies

by platform and is cumulative across all contexts. If an attempt to cache a newly retrieved CRL would exceed its storage limits, the ASA removes the least recently used CRL until more space becomes available.

- Step 4** Check the **Enforce next CRL update** check box to require valid CRLs to have a Next Update value that has not expired. Uncheck the **Enforce next CRL update** check box to let valid CRLs with no Next Update value or a Next Update value that has expired.
- Step 5** Enter the URL for the OCSP server in the **OCSP Options** area. The ASA uses OCSP servers according to the following order:
1. OCSP URL in a match certificate override rule
 2. OCSP URL configured in the selected OCSP Options attribute
 3. AIA field of a user certificate
- Step 6** By default, the **Disable nonce extension** check box is checked, which cryptographically binds requests with responses to avoid replay attacks. This process works by matching the extension in the request to that in the response, ensuring that they are the same. Uncheck the **Disable nonce extension** check box if the OCSP server you are using sends pregenerated responses that do not include this matching nonce extension.
- Step 7** Choose one of the following options in the **Other Options** area:
- Check the **Accept certificates issued by this CA** check box to indicate that the ASA should accept certificates from the specified CA.
 - Check the **Accept certificates issued by the subordinate CAs of this CA** check box to indicate that the ASA should accept certificates from the subordinate CA.
- Step 8** Click **OK** to close this tab, then click **Apply** to save your configuration changes.
-

Add or Import an Identity Certificate

To add or import a new identity certificate configuration, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Remote Access VPN > Certificate Management > Identity Certificates**.
- Step 2** Click **Add**.
- The **Add Identity Certificate** dialog box appears, with the selected trustpoint name displayed at the top.
- Step 3** Click the **Import the identity certificate from a file (PKCS12 format with Certificate(s) + Private Key)** radio button to import an identity certificate from an existing file.
- Step 4** Enter the passphrase used to decrypt the PKCS12 file.
- Step 5** Enter the path name of the file, or click **Browse** to display the **Import ID Certificate File** dialog box. Find the certificate file, then click **Import ID Certificate File**.
- Step 6** Click the **Add a new identity certificate** radio button to add a new identity certificate.
- Step 7** Click **New** to display the **Add Key Pair** dialog box.
- Step 8** Choose the **RSA** or **ECDSA** key type.
- Step 9** Click the **Use default keypair name** radio button to use the default key pair name.

- Step 10** Click the **Enter a new key pair name** radio button, then enter the new name. The ASA supports multiple key pairs.
- Step 11** Choose the modulus size from the drop-down list. If you are not sure of the modulus size, consult Entrust.
- Step 12** Choose the key pair usage by clicking the **General purpose** radio button (default) or **Special** radio button. When you choose the **Special** radio button, the ASA generates two key pairs, one for signature use and one for encryption use. This selection indicates that two certificates are required for the corresponding identity.
- Step 13** Click **Generate Now** to create new key pairs, then click **Show** to display the **Key Pair Details** dialog box, which includes the following display-only information:
- The name of the key pair whose public key is to be certified.
 - The time of day and the date when the key pair is generated.
 - The usage of an RSA key pair.
 - The modulus size (bits) of the key pairs: 512, 768, 1024, and 2048. The default is 1024.
 - The key data, which includes the specific key data in text format.
- Step 14** Click **OK** when you are done.
- Step 15** Choose a certificate subject DN to form the DN in the identity certificate. then click **Select** to display the **Certificate Subject DN** dialog box.
- Step 16** Choose one or more DN attributes that you want to add from the drop-down list, enter a value, then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:
- **Common Name (CN)**
 - **Department (OU)**
 - **Company Name (O)**
 - **Country (C)**
 - **State/Province (ST)**
 - **Location (L)**
 - **E-mail Address (EA)**
- Step 17** Click **OK** when you are done.
- Step 18** Check the **Generate self-signed certificate** check box to create self-signed certificates.
- Step 19** Check the **Act as local certificate authority and issue dynamic certificates to TLS proxy** check box to have the identity certificate act as the local CA.
- Step 20** Click **Advanced** to establish additional identity certificate settings.
- The **Advanced Options** dialog box appears, with the following three tabs: **Certificate Parameters**, **Enrollment Mode**, and **SCEP Challenge Password**.



Note Enrollment mode settings and the SCEP challenge password are not available for self-signed certificates.

- Step 21** Click the **Certificate Parameters** tab, then enter the following information:
- The FQDN, an unambiguous domain name, to indicate the position of the node in the DNS tree hierarchy.

- The e-mail address associated with the identity certificate.
- The ASA IP address on the network in four-part, dotted-decimal notation.
- Check the **Include serial number of the device** check box to add the ASA serial number to the certificate parameters.

Step 22 Click the **Enrollment Mode** tab, then enter the following information:

- Choose the enrollment method by clicking the **Request by manual enrollment** radio button or the **Request from a CA** radio button.
- The enrollment URL of the certificate to be automatically installed through SCEP.
- The maximum number of minutes allowed to retry installing an identity certificate. The default is one minute.
- The maximum number of retries allowed for installing an identity certificate. The default is zero, which indicates an unlimited number of retries within the retry period.

Step 23 Click the **SCEP Challenge Password** tab, then enter the following information:

- The SCEP password
- The SCEP password confirmation

Step 24 Click **OK** when you are done.

Step 25 Click **Add Certificate** in the **Add Identity Certificate** dialog box.

The new identity certificate appears in the Identity Certificates list.

Step 26 Click **Apply** to save the new identity certificate configuration.

Step 27 Click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:

- The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trust points. The values apply to both available and pending status.
- The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Step 28 To remove an identity certificate configuration, select it, then click **Delete**.



Note

After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Configure Certificate Authentication and Renewal Reminders

Store Certificate Chains

When periodic certificate authentication is configured, the ASA stores the certificate chains received from the clients and re-authenticates them periodically. All the PKI checks (including signature validation, revocation checks, validity dates, key usage checks, and so on) that are done during a normal

session establishment are repeated at configured intervals. If a periodic authentication fails, the VPN session is logged off. If it succeeds, the session continues. Periodic Certificate Authentication is not mandatory and is disabled by default.

**Note**

We currently support a feature for IKEv2 where if a session is brought up with a certificate nearing expiration, the session is logged off as soon as the certificate expires. This feature continues as usual.

Restrictions

Periodic certificate authentication will not work with any other Non-VPN sessions/connections. All VPN sessions support the periodic certificate authentication (IKEv1 RA and L2L, IKEv2 RA and L2L including third party standards based, SSL Client and Clientless, L2TP and EZVPN).

Cached CRLs

ASA caches the CRLs received during a certificate validation for a configurable duration and reuses them for validating other certificates within that duration. Caching provides performance benefits by bypassing the step of requesting and verifying the CRLs repeatedly from the server for every connection attempt. A periodic certificate validation is similar to a normal certificate validation.

Local CA

VPN sessions using the certificates issued by the Local CA can also do periodic certificate validation, including revocation checking just like the certificates issued by third-party CAs.

Inactive Sessions

AnyConnect sessions can go into an inactive state after an IKEv2 or SSL tunnel goes down. The session continues to exist to allow session resumption without re-authentication. The periodic authentication continues on the inactive sessions similar to a normal session with the certificate provided during the initial connect. The inactive sessions may be logged off based on the periodic authentication result as well.

Failover

The client certificates are synced to the standby during the session initialization. Active and Standby will maintain and manage their own separate databases for storing these certificates as sessions come up and down. The database is optimized to avoid storing duplicate CA certificates. No duplicate checking is done on ID certificates since the ID certificates are generally unique for each session.

Clustering

Clustering supports L2L VPN only. In addition, the VPN sessions are established only to the master device and are not synced to any other device. If the master goes down, the sessions need to be re-established.

Set a Certificate Expiration Alert (for Identity or CA Certificates)

ASA checks all the CA and ID certificates in the trust points for expiration once every 24 hours. If a certificate is nearing expiration, a syslog will be issued as an alert.

In addition to the renewal reminder, if an already expired certificate is found in the configuration, a syslog is generated once every day to rectify the configuration by either renewing the certificate or removing the expired certificate.

For example, assume that the expiration alerts are configured to begin at 60 days and repeat every 6 days after that. If the ASA is rebooted at 40 days, an alert is sent on that day, and the next alert is sent on the 36th day.

**Note**

Expiration checking is not done on trust pool certificates. The Local CA trust point is treated as a regular trustpoint for expiration checking too.

Procedure

-
- Step 1** Browse to **Configuration > Device Management > Certificate Management > Identity Certificate/CA Certificate**.
- Step 2** Check the **Enable Certificate Expiration Alert** checkbox.
- Step 3** Fill in the desired number of days:
- Send the first alert before—Configure the number of days (1 to 90) before expiration at which the first alert will go out.
 - Repeat the alert for—Configure the alert frequency (1 to 14 days) if the certificate is not renewed. By default, the first alert is sent 60 days prior to expiration and once every week after until the certificate is renewed and removed. Additionally, an alert is sent on the day of the expiration and once every day after that, and irrespective of the alert configuration, an alert is sent every day during the last week of expiration.
-

Export an Identity Certificate

To export an identity certificate, perform the following steps:

Procedure

-
- Step 1** Click **Export** to display the **Export Certificate** dialog box.
- Step 2** Enter the name of the PKCS12 format file to use in exporting the certificate configuration. Alternatively, click **Browse** to display the **Export ID Certificate File** dialog box to find the file to which you want to export the certificate configuration.

- Step 3** Choose the certificate format by clicking the **PKCS12 Format** radio button or the **PEM Format** radio button.
- Step 4** Enter the passphrase used to encrypt the PKCS12 file for export.
- Step 5** Confirm the encryption passphrase.
- Step 6** Click **Export Certificate** to export the certificate configuration.
- An information dialog box appears, informing you that the certificate configuration file has been successfully exported to the location that you specified.
-

Generate a Certificate Signing Request

To generate a certificate signing request to send to Entrust, perform the following steps:

Procedure

- Step 1** Click **Enroll ASA SSL VPN with Entrust** to display the **Generate Certificate Signing Request** dialog box.
- Step 2** Perform the following steps in the **Key Pair** area:
- Choose one of the configured key pairs from the drop-down list.
 - Click **Show** to display the **Key Details** dialog box, which provides information about the selected key pair, including date and time generated, usage (general or special purpose), modulus size, and key data.
 - Click **OK** when you are done.
 - Click **New** to display the **Add Key Pair** dialog box. When you generate the key pair, you can send it to the ASA or save it to a file.
- Step 3** Enter the following information in the **Certificate Subject DN** area:
- The FQDN or IP address of the ASA.
 - The name of the company.
 - The two-letter country code.
- Step 4** Perform the following steps in the **Optional Parameters** area:
- Click **Select** to display the **Additional DN Attributes** dialog box.
 - Choose the attribute to add from the drop-down list, then enter a value.
 - Click **Add** to add each attribute to the attribute table.
 - Click **Delete** to remove an attribute from the attribute table.
 - Click **OK** when you are done.
- The added attributes appear in the **Additional DN Attributes** field.
- Step 5** Enter additional fully qualified domain name information if the CA requires it.
- Step 6** Click **Generate Request** to generate the certificate signing request, which you can then send to Entrust, or save to a file and send later.
- The **Enroll with Entrust** dialog box appears, with the CSR displayed.

- Step 7** Complete the enrollment process by clicking the **request a certificate from Entrust** link. Then copy and paste the CSR provided and submit it through the Entrust web form, provided at <http://www.entrust.net/cisco/>. Alternatively, to enroll at a later time, save the generated CSR to a file, then click the **enroll with Entrust** link on the **Identity Certificates** pane.
- Step 8** Entrust issues a certificate after verifying the authenticity of your request, which may take several days. You then need to install the certificate by selecting the pending request in the **Identity Certificate** pane and clicking **Install**.
- Step 9** Click **Close** to close the **Enroll with Entrust** dialog box.
-

Install Identity Certificates

To install a new identity certificate, perform the following steps:

Procedure

- Step 1** Click **Add** in the **Identity Certificates** pane to display the **Add Identity Certificate** dialog box.
- Step 2** Click the **Add a new identity certificate** radio button.
- Step 3** Change the key pair or create a new key pair. A key pair is required.
- Step 4** Enter the certificate subject DN information, then click **Select** to display the **Certificate Subject DN** dialog box.
- Step 5** Specify all of the subject DN attributes required by the CA involved, then click **OK** to close the **Certificate Subject DN** dialog box.
- Step 6** In the **Add Identity Certificate** dialog box, click **Advanced** to display the **Advanced Options** dialog box.
- Step 7** To continue, see Steps 17 through 23 of the [Add or Import an Identity Certificate, page 20-15](#).
- Step 8** In the **Add Identity Certificate** dialog box, click **Add Certificate**.
The **Identity Certificate Request** dialog box appears.
- Step 9** Enter the CSR file name of type, text, such as `c:\verisign-csr.txt`, then click **OK**.
- Step 10** Send the CSR text file to the CA. Alternatively, you can paste the text file into the CSR enrollment page on the CA website.
- Step 11** When the CA returns the identity certificate to you, go to the **Identity Certificates** pane, select the pending certificate entry, then click **Install**.
The **Install Identity Certificate** dialog box appears.
- Step 12** Choose one of the following options by clicking the applicable radio button:
- **Install from a file.**
Alternatively, click **Browse** to search for the file.
 - **Paste the certificate data in base-64 format.**
Paste the copied certificate data into the area provided.
- Step 13** Click **Install Certificate**.
- Step 14** Click **Apply** to save the newly installed certificate with the ASA configuration.

- Step 15** To show detailed information about the selected identity certificate, click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:
- The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trustpoints. The values apply to both available and pending status.
 - The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
 - The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.
- Step 16** To remove a code signer certificate configuration, select it, and then click **Delete**.

**Note**

After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Import** to reenter all of the certificate configuration information.

Import a Code Signer Certificate

To import a code signer certificate, perform the following steps:

Procedure

- Step 1** In the **Code Signer** pane, click **Import** to display the **Import Certificate** dialog box.
- Step 2** Enter the passphrase used to decrypt the PKCS12-format file.
- Step 3** Enter the name of the file to import, or click **Browse** to display the **Import ID Certificate File** dialog box and search for the file.
- Step 4** Select the file to import and click **Import ID Certificate File**.
The selected certificate file appears in the **Import Certificate** dialog box.
- Step 5** Click **Import Certificate**.
The imported certificate appears in the **Code Signer** pane.
- Step 6** Click **Apply** to save the newly imported code signer certificate configuration.

Export a Code Signer Certificate

To export a code signer certificate, perform the following steps:

Procedure

- Step 1** In the **Code Signer** pane, click **Export** to display the **Export Certificate** dialog box.
- Step 2** Enter the name of the PKCS12 format file to use in exporting the certificate configuration.

- Step 3** In the **Certificate Format** area, to use the public key cryptography standard, which can be base64 encoded or in hexadecimal format, click the **PKCS12 format** radio button. Otherwise, click the **PEM format** radio button.
- Step 4** Click **Browse** to display the **Export ID Certificate File** dialog box to find the file to which you want to export the certificate configuration.
- Step 5** Select the file and click **Export ID Certificate File**.
The selected certificate file appears in the **Export Certificate** dialog box.
- Step 6** Enter the passphrase used to decrypt the PKCS12 format file for export.
- Step 7** Confirm the decryption passphrase.
- Step 8** Click **Export Certificate** to export the certificate configuration.
-

Configure the Local CA Server

To configure a local CA server on the ASA, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Remote Access VPN > Certificate Management > Local Certificate Authority > CA Server**.
- Step 2** To activate the local CA server, check the **Enable Certificate Authority Server** check box. The default setting is disabled (unchecked). After you enable the local CA server, the ASA generates the local CA server certificate, key pair, and necessary database files, then archives the local CA server certificate and key pair in a PKCS12 file.



Note Be sure to review all optional settings carefully before you enable the configured local CA. After you enable it, the certificate issuer name and key size server values cannot be changed.

The self-signed certificate key usage extension enables key encryption, key signature, CRL signature, and certificate signature.

- Step 3** When you enable the local CA for the first time, you must enter and confirm an alphanumeric Enable passphrase, which must have a minimum of seven, alphanumeric characters. The passphrase protects the local CA certificate and the local CA certificate key pair archived in storage, and secures the local CA server from unauthorized or accidental shutdown. The passphrase is required to unlock the PKCS12 archive if the local CA certificate or key pair is lost and must be restored.



Note The Enable passphrase is required to enable the local CA server. Be sure to keep a record of the Enable passphrase in a safe location.

- Step 4** Click **Apply** to save the local CA certificate and key pair, so the configuration is not lost if you reboot the ASA.
- Step 5** To change or reconfigure the local CA after the local CA has been configured for the first time, you must shut down the local CA server on the ASA by unchecking the **Enable Certificate Authority Server** check box. In this state, the configuration and all associated files remain in storage and enrollment is disabled.

After the configured local CA has been enabled, the following two settings are display-only:

- The **Issuer Name** field, which lists the issuer subject name and domain name, and is formed using the username and the subject-name-default DN setting as `cn=FQDN`. The local CA server is the entity that grants the certificate. The default certificate name is provided in the format, `cn=hostname.domainname`.
- The **CA Server Key Size** setting, which is used for the server certificate generated for the local CA server. Key sizes can be 512, 768, 1024, or 2048 bits per key. The default is 1024 bits per key. We recommend that you use a key size of at least 2048.

Step 6 From the drop-down list, choose the client key size of the key pair to be generated for each user certificate issued by the local CA server. Key sizes can be 512, 768, 1024, or 2048 bits per key. The default is 1024 bits per key. We recommend that you use a key size of at least 2048.

Step 7 Enter the CA certificate lifetime value, which specifies the number of days that the CA server certificate is valid. The default is 3650 days (10 years). Make sure that you limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038.

The local CA server automatically generates a replacement CA certificate 30 days before expiration, which enables the replacement certificate to be exported and imported onto any other devices for local CA certificate validation of user certificates that have been issued by the local CA after they have expired.

To notify users of the upcoming expiration, the following syslog message appears in the **Latest ASDM Syslog Messages** pane:

```
%ASA-1-717049: Local CA Server certificate is due to expire in days days and a replacement certificate is available for export.
```



Note When notified of this automatic rollover, the administrator must take action to make sure that the new local CA certificate is imported to all necessary devices before it expires.

Step 8 Enter the client certificate lifetime value, which specifies the number of days that a user certificate issued by the CA server is valid. The default is 365 days (one year). Make sure that you limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038.

Step 9 Set up e-mail access for the local CA server by specifying the following settings in the **SMTP Server & Email Settings** area:

- Enter the SMTP mail server name or IP address. Alternatively, click the ellipses (...) to display the **Browse Server Name/IP Address** dialog box, where you can choose the server name or IP address. Click **OK** when you are done.
- Enter the from address, from which to send e-mail messages to local CA users, in the format "adminname@hostname.com." Automatic e-mail messages carry one-time passwords to newly enrolled users and issue e-mail messages when certificates need to be renewed or updated.
- Enter the subject, which specifies the subject line in all messages that are sent to users by the local CA server. If you do not specify a subject, the default is "Certificate Enrollment Invitation."

Step 10 Click the **More Options** drop-down arrow to configure additional options.

Step 11 Enter the CRL distribution point, which is the CRL location on the ASA. The default location is `http://hostname.domain/+CSCOCA+/asa_ca.crl`.

Step 12 To make the CRL available for HTTP download on a given interface and port, choose a publish-CRL interface from the drop-down list. Then enter the port number, which can be any port number from 1-65535. The default port number is TCP port 80.



Note You cannot rename the CRL; it always has the name, LOCAL-CA-SERVER.crl.

For example, enter the URL, `http://10.10.10.100/user8/my_crl_file`. In this case, only the interface with the specified IP address works and when the request comes in, the ASA matches the path, `/user8/my_crl_file` to the configured URL. When the path matches, the ASA returns the stored CRL file.

Step 13 Enter the CRL lifetime in hours that the CRL is valid. The default for the CA certificate is six hours. The local CA updates and reissues the CRL each time that a user certificate is revoked or unrevoked, but if no revocation changes occur, the CRL is reissued once every CRL lifetime. You can force an immediate CRL update and regeneration by clicking **Request CRL** in the **CA Certificates** pane.

Step 14 Enter the database storage location to specify a storage area for the local CA configuration and data files. The ASA accesses and implements user information, issued certificates, and revocation lists using a local CA database. Alternatively, to specify an external file, enter the path name to the external file or click **Browse** to display the **Database Storage Location** dialog box.

Step 15 Choose the storage location from the list of folders that appears, then click **OK**.



Note Flash memory can store a database with 3500 users or less; a database of more than 3500 users requires external storage.

Step 16 Enter a default subject (DN string) to append to a username on issued certificates. The permitted DN attributes are provided in the following list:

- **CN (Common Name)**
- **SN (Surname)**
- **O (Organization Name)**
- **L (Locality)**
- **C (Country)**
- **OU (Organization Unit)**
- **EA (E-mail Address)**
- **ST (State/Province)**
- **T (Title)**

Step 17 Enter the number of hours for which an enrolled user can retrieve a PKCS12 enrollment file to enroll and retrieve a user certificate. The enrollment period is independent of the one-time password (OTP) expiration period. The default is 24 hours.



Note Certificate enrollment for the local CA is supported only for clientless SSL VPN connections. For this type of connection, communications between the client and the ASA is through a web browser that uses standard HTML.

Step 18 Enter the length of time that a one-time password e-mailed to an enrolling user is valid. The default is 72 hours, then click **Email OTP**.

An **Information** dialog box appears indicating that the OTP was sent to the new user.

Click **Replace OTP** to automatically reissue a new OTP and send an e-mail notice with the new password to an existing or new user.

To view or regenerate the OTP, select a user from the list, then click **View/Regenerate OTP** to display the **View & Regenerate OTP** dialog box.

The current OTP appears.

Click **Regenerate OTP**.

The newly regenerated OTP appears.

Step 19 Click **OK**.

Step 20 Enter the number of days before expiration reminders are e-mailed to users. The default is 14 days.

Step 21 Click **Apply** to save the new or modified CA certificate configuration.

To remove the local CA server from the ASA, click **Delete Certificate Authority Server** to display the **Delete Certificate Authority** dialog box. Click **OK**.



Note After you delete the local CA server, it cannot be restored or recovered. To recreate the deleted CA server configuration, you must reenter all of the CA server configuration information.

.Add a Local CA User

To add a local CA user, perform the following steps:

Procedure

Step 1 To enter a new user into the local CA database, click **Add** to display the **Add User** dialog box.

Step 2 Enter a valid username.

Step 3 Enter an existing valid e-mail address.

Step 4 Enter the subject (DN string). Alternatively, click **Select** to display the **Certificate Subject DN** dialog box.

Step 5 Choose one or more DN attributes that you want to add from the drop-down list, enter a value, and then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:

- **Common Name (CN)**
- **Department (OU)**
- **Company Name (O)**
- **Country (C)**
- **State/Province (ST)**
- **Location (L)**
- **E-mail Address (EA)**

Step 6 Click **OK** when you are done.

Step 7 Check the **Allow enrollment** check box to enroll the user, then click **Add User**.

The new user appears in the **Manage User Database** pane.

Edit a Local CA User

To modify information about an existing local CA user in the database, perform the following steps:

-
- Step 1** Select the specific user and click **Edit** to display the **Edit User** dialog box.
- Step 2** Enter a valid username.
- Step 3** Enter an existing valid e-mail address.
- Step 4** Enter the subject (DN string). Alternatively, click **Select** to display the **Certificate Subject DN** dialog box.
- Step 5** Choose one or more DN attributes that you want to change from the drop-down list, enter a value, and then click **Add** or **Delete**.
- Step 6** Click **OK** when you are done.

To remove the user from the database and any certificates issued to that user from the local CA database, select the user, then click **Delete**.



Note A deleted user cannot be restored. To recreate the deleted user record, click **Add** to reenter all of the user information.

- Step 7** Check the **Allow enrollment** check box to re-enroll the user, then click **Edit User**.



Note If the user is already enrolled, an error message appears.

The updated user details appear in the **Manage User Database** pane.

Manage User Certificates

To change the certificate status, perform the following steps:

Procedure

-
- Step 1** Select specific certificates by username or by certificate serial number in the **Manage User Certificates** pane.
- Step 2** Choose one of the following options:
- Click **Revoke** to remove user access if the user certificate lifetime period runs out. The local CA also marks the certificate as revoked in the certificate database, automatically updates the information, and reissues the CRL.
 - Select a revoked certificate and click **Unrevoke** to restore access. The local CA also marks the certificate as unrevoked in the certificate database, automatically updates the certificate information, and reissues an updated CRL.
- Step 3** Click **Apply** when you are done to save your changes.
-

Monitoring Digital Certificates

See the following screens for monitoring digital certificate status:

- **Monitoring > Properties > CRL**

This pane shows CRL details.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for Certificate Management

Table 20-1 History for Certificate Management

Feature Name	Platform Releases	Description
Certificate management	7.0(1)	Digital certificates (including CA certificates, identity certificates, and code signer certificates) provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security. We introduced the following screens: Configuration > Remote Access VPN > Certificate Management Configuration > Site-to-Site VPN > Certificate Management. We introduced or modified the following screens: Configuration > Firewall > Advanced > Certificate Management > CA Certificates Configuration > Device Management > Certificate Management > CA Certificates.
Certificate management	7.2(1)	
Certificate management	8.0(2)	
SCEP proxy	8.4(1)	We introduced this feature, which provides secure deployment of device certificates from third-party CAs.



PART 5

IP Routing



Routing Overview

This chapter describes underlying concepts of how routing behaves within the Cisco ASA, and the routing protocols that are supported. Routing is the act of moving information across a network from a source to a destination. Along the way, at least one intermediate node is typically encountered. Routing involves two basic activities: determining optimal routing paths and transporting packets through a network.

- [Path Determination, page 21-1](#)
- [Supported Route Types, page 21-2](#)
- [How Routing Behaves Within the ASA, page 21-3](#)
- [Supported Internet Protocols for Routing, page 21-5](#)
- [Routing Table, page 21-6](#)
- [Disable Proxy ARP Requests, page 21-10](#)
- [Display the Routing Table, page 21-11](#)

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 21-2](#)
- [Single-Path Versus Multipath, page 21-3](#)
- [Flat Versus Hierarchical, page 21-3](#)
- [Link-State Versus Distance Vector, page 21-3](#)

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.

- [Egress Interface Selection Process, page 21-4](#)
- [Next Hop Selection Process, page 21-4](#)
- [ECMP Routing, page 21-4](#)

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.

ECMP Routing

The ASA supports Equal-Cost Multi-Path (ECMP) routing.

Without zones, you can have up to 3 equal cost static or dynamic routes per interface. For example, you can configure three default routes on the outside interface that specify different gateways:

```
route outside 0 0 10.1.1.2
route outside 0 0 10.1.1.3
route outside 0 0 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

ECMP is not supported across multiple interfaces, so you cannot define a route to the same destination on a different interface. The following route is disallowed when configured with any of the routes above:

```
route outside2 0 0 10.2.1.1
```

With zones, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within a zone. For example, you can configure three default routes across three interfaces in the zone:

```
route outside1 0 0 10.1.1.2
route outside2 0 0 10.2.1.2
route outside3 0 0 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The ASA load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the ASA seamlessly moves the flow to a different route.

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 [Configure EIGRP, page 27-3](#).

- 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 [Configure OSPFv2, page 26-6](#).

- 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 legacy feature guide.

- Border Gateway Protocol (BGP)

BGP is an interautonomous system routing protocol. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

For more information about configuring BGP, see [Configure BGP, page 25-4](#)

Routing Table

- [How the Routing Table Is Populated, page 21-6](#)
- [How Forwarding Decisions Are Made, page 21-8](#)
- [Dynamic Routing and Failover, page 21-8](#)
- [Dynamic Routing and Clustering, page 21-9](#)
- [Dynamic Routing in Multiple Context Mode, page 21-9](#)

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, OSPF and BGP 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 21-1 shows the default administrative distance values for the routing protocols supported by the ASA.

Table 21-1 **Default Administrative Distance for Supported Routing Protocols**

Route Source	Default Administrative Distance
Connected interface	0
Static route	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
RIP	120
EIGRP external route	170
Internal BGP	200
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the ASA receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the ASA chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the ASA would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you use the **distance-ospf** command to change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the ASA on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The EIGRP, OSPF, RIP and BGP 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

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.

Dynamic Routing and Clustering

Dynamic routing is fully integrated in a cluster, and routes are shared across units. 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 Spanned EtherChannel 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 Spanned EtherChannel interfaces are replicated.

When new routing entries are learned through Spanned EtherChannel 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.

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.

If you configure Individual interfaces, you must also configure the router-id pool setting.

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* 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.

Disable 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.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Routing > Proxy ARP/Neighbor Discovery**.
- The Interface field lists the interface names. The **Enabled** field shows whether or not proxy ARP/Neighbor Discovery is enabled (Yes) or disabled (No) for NAT global addresses.
- Step 2** To enable proxy ARP/Neighbor Discovery for the selected interface, click **Enable**. By default, proxy ARP/Neighbor discovery is enabled for all interfaces.
- Step 3** To disable proxy ARP/Neighbor Discovery for the selected interface, click **Disable**.
- Step 4** Click **Apply** to save your settings to the running configuration.
-

Display the Routing Table

Procedure

-
- Step 1** To show all routes in ASDM that are in the routing table, choose **Monitoring > Routing > Routes**. In this pane, each row represents one route.
-



Static and Default Routes

This chapter describes how to configure static and default routes on the Cisco ASA.

- [About Static and Default Routes, page 22-1](#)
- [Guidelines for Static and Default Routes, page 22-3](#)
- [Configure Default and Static Routes, page 22-3](#)
- [Monitoring a Static or Default Route, page 22-7](#)
- [Examples for Static or Default Routes, page 22-7](#)
- [History for Static and Default Routes, page 22-7](#)

About Static and Default Routes

To route traffic to a nonconnected host or network, you must define a route to the host or network, either using static or dynamic routing. Generally, you must configure at least one static route: a default route for all traffic that is not routed by other means to a default network gateway, typically the next hop router.

- [Default Route, page 22-1](#)
- [Static Routes, page 22-2](#)
- [Route to null0 Interface to “Black Hole” Unwanted Traffic, page 22-2](#)
- [Route Priorities, page 22-2](#)
- [Transparent Firewall Mode Routes, page 22-2](#)
- [Equal-Cost Multi-Path Routes, page 22-2](#)
- [Static Route Tracking, page 22-3](#)

Default Route

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. 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.

Static Routes

You might want to use static routes in single context mode in the following cases:

- Your networks use a different router discovery protocol from BGP, 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.
- In some cases, a default route is not enough. 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.

Route to null0 Interface to “Black Hole” Unwanted Traffic

Access rules let you filter packets based on the information contained in their headers. A static route to the null0 interface is a complementary solution to access rules. You can use a null0 route to forward unwanted or undesirable traffic into a “black hole” so the traffic is dropped.

Static null0 routes have a favorable performance profile. You can also use static null0 routes to prevent routing loops. BGP can leverage the static null0 route for Remotely Triggered Black Hole routing.

Route Priorities

- Routes that identify a specific destination take precedence over the default route.
- When multiple routes exist to the same destination (either static or dynamic), then the administrative distance for the route determines priority. Static routes are set to 1, so they typically are the highest priority routes.
- When you have multiple static routes to the same destination with the same administrative distance, see [ECMP Routing, page 21-4](#).
- For traffic emerging from a tunnel with the Tunneled option, this route overrides any other configured or learned default routes.

Transparent Firewall Mode Routes

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.

Equal-Cost Multi-Path Routes

The ASA supports Equal-Cost Multi-Path (ECMP) routing. See [ECMP Routing, page 21-4](#) for more information.

Static Route Tracking

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.

Guidelines for Static and Default Routes

Firewall Mode

Static route tracking is not supported in transparent firewall mode.

IPv6

- IPv6 static routes are not supported in transparent mode in ASDM.
- Static route tracking is not supported for IPv6.

Clustering

In clustering, static route monitoring is only supported on the master unit.

Configure Default and Static Routes

At a minimum, you should configure a default route. You may need to configure static routes as well.

- [Configure a Default Route, page 22-3](#)
- [Configure a Static Route, page 22-4](#)
- [Track Static Routes, page 22-5](#)

Configure a Default Route

A default route is simply a static route with 0.0.0.0/0 as the destination IP address.

Before You Begin

See the following guidelines for the Tunneled option:

- Do not enable unicast RPF 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 inspection 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.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Routing > Static Routes**, and click **Add**.
- Step 2** Choose the **IP Address Type**, **IPv4** or **IPv6**.
- Step 3** Choose the **Interface** through which you want to send the traffic.
- Step 4** For the **Network**, type **any4** or **any6**, depending on the type.
- Step 5** Enter the **Gateway IP** where you want to send the traffic.
- Step 6** Set the **Metric** to set the administrative distance for the route.
- The default is **1**. 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.
- Step 7** (Optional) In the **Options** area, set the following:
- **Tunneled**—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.
 - **Tracked**—(IPv4 only) For information about tracking routes, see [Track Static Routes, page 22-5](#).
- Step 8** Click **OK**.
-

Configure a Static Route

A static route defines where to send traffic for specific destination networks.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Routing > Static Routes**, and click **Add**.
- Step 2** Choose the **IP Address Type**, **IPv4** or **IPv6**.
- Step 3** Choose the **Interface** through which you want to send the traffic. To “black hole” unwanted traffic, choose the **Null0** interface.
- Step 4** For the **Network**, enter the destination network for which you want to route traffic.
- Step 5** Enter the **Gateway IP** where you want to send the traffic.
- Step 6** Set the **Metric** to set the administrative distance for the route.
- The default is **1**. 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.

Step 7 (Optional) In the **Options** area, set the following:

- **Tunneled**—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.
- **Tracked**—(IPv4 only) For information about tracking routes, see [Track Static Routes, page 22-5](#).

Step 8 Click **OK**.

Track Static Routes

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.

- [About Static Route Tracking, page 22-5](#)
- [Configure Static Route Tracking, page 22-5](#)

About Static Route Tracking

The ASA implements static route tracking by associating a static route with a monitoring target host on the destination network that the ASA monitors using ICMP echo requests. If an echo reply is not received within a specified time period, the host is considered down, and the associated route is removed from the routing table. An untracked backup route with a higher metric 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 PC 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.

Configure Static Route Tracking

To configure static route tracking, complete the following steps.

Before You Begin

Static route tracking is supported for:

- IPv4 traffic.
- Routed firewall mode.

Procedure

Step 9

- Step 1** Choose **Configuration > Device Setup > Routing > Static Routes** and add or edit a static route according to [Configure a Static Route, page 22-4](#).
- Step 2** Click the **Tracked** radio button in the **Options** area.
- Step 3** In the **Track ID** field, enter a unique identifier for the route tracking process.
- Step 4** In the **Track IP Address/DNS Name** field, enter the IP address or hostname of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available from that interface.
- Step 5** In the **SLA ID** field, enter a unique identifier for the SLA monitoring process.
- Step 6** (Optional) Click **Monitoring Options**.

The **Route Monitoring Options** dialog box appears. From here, change the following tracking object monitoring properties:

- **Frequency**—Sets how often, in seconds, the ASA should test for the presence of the tracking target. Valid values range from 1 to 604800 seconds. The default value is 60 seconds.
- **Threshold**—Sets the amount of time, in milliseconds, that indicates an over-threshold event. This value cannot be more than the timeout value.
- **Timeout**—Sets the amount of time, in milliseconds, that the route monitoring operation should wait for a response from the request packets. Valid values range from 0 to 604800000 milliseconds. The default value is 5000 milliseconds.
- **Data Size**—Sets the size of data payload to use in the echo request packets. The default value is 28. Valid values range from 0 to 16384.



Note This setting specifies the size of the payload only; it does not specify the size of the entire packet.

- **ToS**—Sets a value for the type of service byte in the IP header of the echo request. Valid values are from 0 to 255. The default value is 0.
- **Number of Packets**—Sets the number of echo requests to send for each test. Valid values range from 1 to 100. The default value is 1.

Click **OK**.

- Step 7** Click **OK** to save your route, and then click **Apply**.

The monitoring process begins as soon as you apply the tracked route.

- Step 8** Create an untracked backup route.

The backup route is a static route to the same destination as the tracked route, but through a different interface or gateway. You must assign this route a higher administrative distance (metric) than your tracked route.

Monitoring a Static or Default Route

- **Monitoring > Routing > Routes.**

In the **Routes** pane, each row represents one route. You can filter by IPv4 connections, IPv6 connections, or both. The routing information includes the protocol, the route type, the destination IP address, the netmask or prefix length, the gateway IP address, the interface through which the route is connected, and the administrative distance.

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 dmz interface, and adds a default route for tunneled traffic and one for regular traffic.

```
route inside 10.1.1.0 255.255.255.0 10.1.2.45
route dmz 10.10.10.0 255.255.255.0 192.168.2.1
route dmz 10.10.10.0 255.255.255.0 192.168.2.2
route dmz 10.10.10.0 255.255.255.0 192.168.2.3
route outside 0 0 209.165.201.1
route inside 0 0 10.1.2.45 tunneled
```

History for Static and Default Routes

ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 22-1 Feature History for Static and Default Routes

Feature Name	Platform Releases	Feature Information
Static Route Tracking	7.2(1)	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. We introduced or modified the following screens: Configuration > Device Setup > Routing > Static Routes > Add Static Route Configuration > Device Setup > Routing > Static Routes > Add Static Route > Route Monitoring Options
Static null0 route to “black hole” traffic	9.2(1)	Sending traffic to a null0 interface results in dropping the packets destined to the specified network. This feature is useful in configuring Remotely Triggered Black Hole (RTBH) for BGP. We modified the following screen: Configuration > Device Setup > Routing > Static Routes > Add Static Route



Policy Based Routing

This chapter describes how to configure the Cisco ASA to support policy based routing (PBR). The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- [About Policy Based Routing, page 23-1](#)
- [Guidelines for Policy Based Routing, page 23-3](#)
- [Configure Policy Based Routing, page 23-4](#)
- [History for Policy Based Routing, page 23-6](#)

About Policy Based Routing

Typical routing systems and protocols are based on routing the traffic based on the destination of the traffic. Destination based routing systems make it hard to change the routing behavior of specific traffic. With Policy Based Routing (PBR), you can define the routing behavior based on a number of different criteria other than destination network. These criteria include source or destination network, source or destination address, source or destination port, protocol, packet size, and packet classification among several others.

PBR has the ability to implement Quality of Service (QoS) by classifying and marking traffic at the network edge and then using PBR throughout the network to route marked traffic along a specific path.

This permits routing of packets originating from different sources to different networks even when the destinations are the same and can be useful when interconnecting several private networks.

Why Use Policy Based Routing?

Consider a company that has two links between locations, one a high bandwidth, low delay expensive link and the other a low bandwidth, higher delay lower expense link. While using traditional routing protocols the higher bandwidth link would get most if not all of the traffic sent across it based on the metric savings obtained by the bandwidth and/or delay (using EIGRP or OSPF) characteristics of the link. PBR allows you to route higher priority traffic over the high bandwidth/low delay link and sends all other traffic over the low bandwidth/high delay link.


Some applications of policy based routing are listed below:

- Equal-Access and Source-Sensitive routing
- Quality of Service
- Cost Saving

- Load Sharing


Equal-Access and Source-Sensitive Routing

In this topology, traffic from HR network & Mgmt network can be configured to go through ISP1 and traffic from Eng network can be configured to go through ISP2. Thus, policy based routing enables the network administrators to provide equal-access and source-sensitive routing as shown below.




Quality of Service

By tagging packets with policy based routing, network administrators can classify the network traffic at the perimeter of the network for various classes of service and then implementing those classes of service in the core of the network using priority, custom or weighted fair queuing (as shown in the figure below). This setup improves network performance by eliminating the need to classify the traffic explicitly at each WAN interface in the core of backbone network.



Cost Saving

An organization can direct the bulk traffic associated with a specific activity to use a higher-bandwidth high-cost link for a short time and continues basic connectivity over a lower-bandwidth low-cost link for interactive traffic by defining the topology as show below.



Load Sharing

In addition to the dynamic load-sharing capabilities offered by ECMP load balancing, network administrators can now implement policies to distribute traffic among multiple paths based on the traffic characteristics.

As an example, in the topology depicted in the Equal-Access Source Sensitive Routing scenario, an administrator can configure policy based routing to load share the traffic from HR network through ISP1 and traffic from Eng network through ISP2.

Implementation of PBR

The ASA uses ACLs to match traffic and then perform routing actions on the traffic. Specifically, you configure a route map that specifies an ACL for matching, and then you specify one or more actions for that traffic. Finally, you associate the route map with an interface where you want to apply PBR on all incoming traffic

Guidelines for Policy Based Routing

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Clustering

Clustering is not supported.

Additional Guidelines

- All existing route map related configuration restrictions and limitations will be carried forward.

Configure Policy Based Routing

A route map is comprised of one or more route-map statements. Each statement has a sequence number, as well as a permit or deny clause. Each route-map statement contains match and set commands. The match command denotes the match criteria to be applied on the packet. The set command denotes the action to be taken on the packet.

- When a route map is configured with both IPv4 and IPv6 match/set clauses or when a unified ACL matching IPv4 and IPv6 traffic is used, the set actions will be applied based on destination IP version.
- When multiple next-hops or interfaces are configured as set action, all these options are evaluated one after the other until a valid usable option is found. No load balancing will be done among the configured multiple options.
- When a route map with match ACLs containing IPv6 addresses is attached to an interface, all IPv6 related ACLs will be discarded with a warning.
- Verify-availability option uses track object configuration to track and verify the status of configured next-hop. Since track object configuration is not supported in multi-mode, we will not support verify-availability option too.

Procedure

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

The Add Route Map or Edit Route Map dialog box appears.

Step 3 Enter the route map name and sequence number. The route map name is the name that you assign to a particular route. The sequence number is the order in which you add or delete the route map entries into the ASA.



Note If you are editing an existing route map, the fields for Route Map name and sequence number are already filled in.

Step 4 To reject route matches from redistribution, click **Deny**. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. To allow route matches for redistribution, click **Permit**. If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed.

In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.


Step 5 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- In the IPv4 and IPv6 sections, do one or more of the following, choose the ACL or prefix list that you want.

Step 6 Click the **Policy Based Routing** tab to define policy for traffic flows, and lessening reliance on routes derived from routing protocols. PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.

- Check the **Set default next-hop IP address** check box to indicate where to output packets that pass a match clause of a route map for policy routing. In the **IPv4 Address** enter the destination address.
- Check the **Recursively find and set next-hop IP address** check box and specify an IP address in the **IPv4 Address** field. In this case, the next-hop IP address need not be on a directly connected subnet.
- Check the **Verify the availability of the next hop IPv4 address** check box to verify if the next hops of a route map are Cisco Discovery Protocol (CDP) neighbors before policy routing to those next hops. Here a tracking object is created to track the reachability of a next-hop using ICMP echo/reply and mark the next-hop as unreachable when it does not reply to the ICMP echo requests.
 - In the **IPv4 Address** field, enter the next hop IP address.
 - In the **Sequence number** field and the **Track** field, enter valid values that verify the reachability of the tracked object.
- Check the **Set interfaces** check box and select a destination interface from the drop-down list.
- Check the **Set null0 interface as the default interface** check box, if there is a need to completely black hole or drop some traffic.
- Check the **Set do-not-fragment bit to either 1 or 0** and then select the appropriate radio button.
- Check the **Set differential service code point (DSCP) value in QoS bits for IPv4 packets** check box and select a value from the **IPv4 DSCP value** drop-down list.

Step 7 Click **OK**.



History for Policy Based Routing

Table 23-1 History for Route Maps

Feature Name	Platform Releases	Feature Information
Policy based routing	9.4(1)	Policy Based Routing (PBR) is a mechanism by which traffic is routed through specific paths with a specified QoS using ACLs. ACLs let traffic be classified based on the content of the packet's Layer 3 and Layer 4 headers. This solution lets administrators provide QoS to differentiated traffic, distribute interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths, and allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections. We updated the following screens: Configuration > Device Setup > Routing > Route Maps >Policy Based Routing, Configuration > Device Setup > Routing > Interface Settings > Interfaces.
IPv6 support for Prefix Rule	9.3.2	We introduced this feature. We updated the following screen: Configuration > Device Setup > Routing > Prefix Rules > Add prefix Rule
Policy Based Routing for Route Maps	9.4.1	Policy Based Routing (PBR) is a mechanism by which traffic is routed through specific paths with a specified QoS using ACLs. ACLs let traffic be classified based on the content of the packet's Layer 3 and Layer 4 headers. This solution lets administrators provide QoS to differentiated traffic, distribute interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths, and allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections. We updated the following screens: Configuration > Device Setup > Routing > Route Maps >Policy Based Routing, Configuration > Device Setup > Routing > Interface Settings > Interfaces.



Route Maps

This chapter describes how to configure and customize route-maps, for Cisco ASA.

- [About Route Maps, page 24-1](#)
- [Guidelines for Route Maps, page 24-4](#)
- [Define a Route Map, page 24-4](#)
- [Customize a Route Map, page 24-6](#)
- [Configuration Example for Route Maps, page 24-9](#)
- [Feature History for Route Maps, page 24-10](#)

About Route Maps

Route maps are used when redistributing routes into an OSPF, RIP, EIGRP or BGP 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 Cisco ASDM, this capability for redistribution can be found when you add or edit a new route map (see [Define a Route Map, page 24-4](#)). 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.

- [Permit and Deny Clauses, page 24-2](#)
- [Match and Set Clause Values, page 24-2](#)
- [BGP Match and BGP Set Clauses, page 24-3](#)
- [Guidelines for Route Maps, page 24-4](#)

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).

BGP Match and BGP Set Clauses

In addition to the match and set values described above, BGP provides additional match and set capabilities to route maps.

The following new route-map match clauses are now supported with BGP:

- match as-path
- match community
- match policy-list
- match tag

The following new route-map set clauses are now supported with BGP:

- set as-path
- set automatic-tag
- set community
- set local-preference
- set origin
- set weight

For each BGP route that is being redistributed, the ASA first evaluates the BGP match criteria of a clause in the route map. If the BGP match criteria succeeds, 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 BGP Set Clause 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), as set from the BGP Match Clause tab in ASDM, match the route or until the end of the route map is reached.

Guidelines for Route Maps

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Additional Guidelines

Route maps do not support ACLs that include a user, user group, or fully qualified domain name objects.

Define 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. In ASDM, you can define a route map by adding, editing, or deleting a route map name, sequence number, or redistribution.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.
 - Step 2** Click **Add**.
The Add Route Map or Edit Route Map dialog box appears.
 - Step 3** Enter the route map name and sequence number. The route map name is the name that you assign to a particular route. The sequence number is the order in which you add or delete the route map entries into the ASA.



Note If you are editing an existing route map, the fields for Route Map name and sequence number are already filled in.

- Step 4** To reject route matches from redistribution, click **Deny**. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. To allow route matches for redistribution, click **Permit**. If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed.

In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.
- Step 5** Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
 - Enter the interface name in the Interface field, or click the ellipses to display the Browse Interface dialog box.
 - Choose one or more interfaces, click **Interface**, then click **OK**.
- In the IPv4 and IPv6 sections, do one or more of the following:
 - Check the **Match Address** check box to enable or disable the Match address of a route or match packet.
 - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.
 - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
 - Choose Access List to Prefix List from the drop-down list to match the IP address.
 - According to the previous selection, click the ellipses to display the Browse Access List or Browse Prefix List dialog box.
 - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
 - In the Metric Value field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.
- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Step 6 Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:

- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and type the value in the Value field.
- Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.

Step 7 Click the **BGP Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match AS path access lists** check box to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list.
- Check the **Match Community** check box to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Any route that does not match atleast one Match community will not be advertised for outbound route maps.
 - Check the **Match the specified community exactly** check box to enable matching the BGP community exactly with the specified community.
- Check the **Match Policy list** check box to configure a route map to evaluate and process a BGP policy. If you specify more than one policy list, then the route can process either policy list.

- Step 8** Click the **BGP Set Clause** tab to modify the following information, which will be redistributed to the BGP protocol:
- Check the **Set AS Path** check box to modify an autonomous system path for BGP routes.
 - Check the **Prepend AS path** check box to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS numbers.
 - Check the **Prepend Last AS to the AS Path** check box to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10.
 - Check the **Convert route tag into AS Path** check box to convert the tag of a route into an autonomous system path.
 - Check the **Set Community** check box to set the BGP communities attributes.
 - Click **Specify Community** to enter a community number, if applicable. Valid values are from 1 to 4294967200, internet, no-advertise and no-export.
 - Check **Add to the existing communities** to add the community to the already existing communities.
 - Click **None** to remove the community attribute from the prefixes that pass the route map.
 - Check the **Set local preference** check box to specify a preference value for the autonomous system path.
 - Check the **Set weight** check box to specify the BGP weight for the routing table. Enter a value between 0 and 65535.
 - Check the **Set origin** check box to specify the BGP origin code. Valid values are Local IGP and Incomplete.
 - Check the **Set next hop** check box to specify the output address of packets that fulfill the match clause of a route map.
 - Click **Specify IP address** to enter the IP address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IP address then the packets can output at either IP address.
 - Click **Use peer address** to set the next hop to be the BGP peer address.

Step 9 Click **OK**.

Customize a Route Map

This section describes how to customize the route map.

- [Define a Route to Match a Specific Destination Address, page 24-7](#)
- [Configure Prefix Rules, page 24-7](#)
- [Configure Prefix Lists, page 24-8](#)
- [Configure the Metric Values for a Route Action, page 24-9](#)

Define a Route to Match a Specific Destination Address

Procedure

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.

Step 2 Click **Add**.

The Add Route Map dialog box appears. From this dialog box you can assign or choose the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add the entries.

Step 3 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
 - Enter the interface name in the Interface field, or click the ellipses to display the Browse Interface dialog box.
 - Choose the interface type (**inside** or **outside**), click **Selected Interface**, then click **OK**.
 - Check the **Match IP Address** check box to enable or disable the Match address of a route or match packet.
 - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.
 - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
 - Choose Access List to Prefix List from the drop-down list to match the IP address.
 - According to the previous selection, click the ellipses to display the Browse Access List or Browse Prefix List dialog box.
 - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
 - In the Metric Value field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.
- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Configure Prefix Rules



Note

You must configure a prefix list before you may configure a prefix rule.

To configure prefix rules, perform the following steps:

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > Prefix Rules**.
- Step 2** Click **Add** and choose **Add Prefix Rule**.
- The Add Prefix Rule dialog box appears. From this dialog box, you can add a sequence number, select an IP version- IPv4 or IPv6, specify a prefix for the network, its redistribution access (that is, permit or deny) and the minimum and maximum prefix length.
- Step 3** Enter an optional sequence number or accept the default value.
- Step 4** Specify the prefix number in the format of IP address/mask length.
- Step 5** Click the **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 6** Enter the optional minimum and maximum prefix lengths.
- Step 7** Click **OK** when you are done.
- The new or revised prefix rule appears in the list.
- Step 8** Check the **Enable Prefix list sequence numbering** check box if you want to use automatically generated sequence numbers.
- Step 9** Click **Apply** to save your changes.
-

Configure Prefix Lists

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

To add prefix lists, perform the following steps:

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > Prefix Rules**.
- Step 2** Click **Add** and choose **Add Prefix List**.
- The Add Prefix List dialog box appears.
- Step 3** Enter the prefix name and description, then click **OK**.
-

Configure the Metric Values for a Route Action

To configure the metric value for a route action, perform the following steps:

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.
- Step 2** Click **Add**.
- The Add Route Map or Edit Route Map dialog box appears. From this dialog box, you can assign or select the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add route map entries.
- Step 3** Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:
- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and enter the value in the Value field.
 - Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.
-

Configuration Example for Route Maps

The following example shows how to redistribute routes with a hop count equal to 1 into OSPF.

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > Route Maps**.
- Step 2** Click **Add**.
- Step 3** Enter **1-to-2** in the Route Map Name field.
- Step 4** Enter the routing sequence number in the Sequence Number field.
- Step 5** Click the **Permit** radio button.
- By default this tab is on top.
- Step 6** Click the **Match Clause** tab.
- Step 7** Check the **Match Metric of Route** check box and type **1** for the metric value.
- Step 8** Click the **Set Clause** tab.
- Step 9** Check the **Set Metric Value** check box, and type **5** for the metric value.
- Step 10** Check the **Set Metric-Type** check box, and choose **Type-1**.
-

Feature History for Route Maps

Table 24-1 Feature History for Route Maps

Feature Name	Platform Releases	Feature Information
Route maps	7.0(1)	We introduced this feature. We introduced the following screen: Configuration > Device Setup > Routing > Route Maps.
Enhanced support for static and dynamic route maps	8.0(2)	Enhanced support for dynamic and static route maps was added.
Dynamic Routing in Multiple Context Mode	9.0(1)	Route maps are supported in multiple context mode.
Support for BGP	9.2(1)	We introduced this feature. We updated the following screen: Configuration > Device Setup > Routing > Route Maps with 2 additional tabs BGP match clause and BGP set clause.
IPv6 support for Prefix Rule	9.3.2	We introduced this feature. We updated the following screen: Configuration > Device Setup > Routing > Prefix Rules > Add prefix Rule



BGP

This chapter describes how to configure the Cisco ASA to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- [About BGP, page 25-1](#)
- [Guidelines for BGP, page 25-3](#)
- [Configure BGP, page 25-4](#)
- [Monitoring BGP, page 25-22](#)
- [History for BGP, page 25-22](#)

About BGP

BGP is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

- [When to Use BGP, page 25-1](#)
- [Routing Table Changes, page 25-2](#)

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

BGP can now be used for carrying routing information for IPv6 prefix over IPv6 networks.



Note

When a BGPv6 ASA joins the cluster, it generates a soft traceback when logging level 7 is enabled.

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- **Weight** -- This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
- **Local preference** -- The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- **Multi-exit discriminator** -- The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- **Origin** -- The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
 - **IGP**- The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
 - **EGP**-The route is learned via the Exterior Border Gateway Protocol (EBGP).
 - **Incomplete**- The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- **AS_path** -- When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- **Next hop** -- The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS.
- **Community** -- The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
 - **no-export**- Do not advertise this route to EBGP peers.
 - **no-advertise**- Do not advertise this route to any peer.
 - **internet**- Advertise this route to the Internet community; all routers in the network belong to it.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

Guidelines for BGP

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Does not support transparent firewall mode. BGP is supported only in router mode.

Failover Guidelines

Supports Stateful Failover in single and multiple context mode.

**Note**

When clustering is enabled, failover is not supported.

Clustering Guidelines

BGP is supported only in L2 (EtherChannel Type) and L3 (Individual Interface Type) clustering modes.

**Note**

When you delete and reapply the BGP configuration in the user context allow a delay of 60 seconds, to enable the slave/ standby ASA unit to sync.

IPv6 Guidelines

Supports IPv6. Graceful restart is not supported for IPv6 address family.

Configure BGP

This section describes how to enable and configure the BGP process on your system.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP**.
 - Step 2** Enable the BGP routing process by checking the **Enable BGP routing** check box on the **General** tab. See [Enable BGP, page 25-4](#).
 - Step 3** Define the configuration related to the best path selection process for BGP routing on the **BGP > Best Path** tab. See [Define the Best Path for a BGP Routing Process, page 25-5](#).
 - Step 4** Configure the Policy Lists for BGP routing on the **BGP > Policy Lists** tab. See [Configure Policy Lists, page 25-6](#).
 - Step 5** Configure the AS Path Filters for BGP routing on the **BGP > AS Path Filters** tab. See [Configure AS Path Filters, page 25-7](#).
 - Step 6** Configure Community Rules for BGP routing on the **BGP > Community Rules** tab. See [Configure Community Rules, page 25-8](#).
 - Step 7** Configure IPv4 Address Family Settings on the **BGP > IPv4 Family** tab. See [Configure IPv4 Address Family Settings, page 25-9](#).
-

Enable BGP

This section describes the steps required to enable BGP routing, establish a BGP routing process and configure general BGP parameters.

Procedure

-
- Step 1** For single-mode, in ASDM, choose **Configuration > Device Setup > Routing > BGP > General**.



Note For multi-mode, in ASDM choose **Configuration > Context Management > BGP**. After enabling BGP, switch to a security context and enable BGP by choosing **Configuration > Device Setup > Routing > BGP > General**.

The **General** pane appears.

- Step 2** Check the **Enable BGP Routing** check box.
- Step 3** In the **AS Number** field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to XX.YY.

- Step 4** (Optional) Check the **Limit the number of AS numbers in the AS_PATH attribute of received routes** check box to restrict the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
- Step 5** (Optional) Check the **Log neighbor changes** check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.
- Step 6** (Optional) Check the **Use TCP path MTU discovery** check box to use the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
- Step 7** (Optional) Check the **Enable fast external failover** check box to reset the external BGP session immediately upon link failure.
- Step 8** (Optional) Check the **Enforce that first AS is peer's AS for EBGP routes** check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
- Step 9** (Optional) Check the **Use dot notation for AS numbers** check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65535 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
- Step 10** Specify the timer information in the **Neighbor timers** area:
- a. In the **Keepalive interval** field, enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.
 - b. In the **Hold Time** field, enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. The default value is 180 seconds.
 - c. (Optional) In the **Min. Hold Time** field, enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.
- Step 11** (Optional) In the **Non Stop Forwarding** section do the following:
- a. Check the **Enable Graceful Restart** check box to enable ASA peers to avoid a routing flap following a switchover.
 - b. In the **Restart Time** field, enter the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
 - c. In the **Stale Path Time** field, enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
- Step 12** Click **OK**.
- Step 13** Click **Apply**.
-

Define the Best Path for a BGP Routing Process

This section describes the steps required to configure the BGP best path. For more information on the best path, see [BGP Path Selection, page 25-3](#).

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Best Path**.
- The **Best Path configuration** pane appears.
- Step 2** In the **Default Local Preference** field, specify a value between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.
- Step 3** Check the **Allow comparing MED from different neighbors** check box to allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
- Step 4** Check the **Compare router-id for identical EBGp paths** check box to compare similar paths received from external BGP peers during the best path selection process and switch the best path to the route with the lowest router ID.
- Step 5** Check the **Pick the best MED path among paths advertised from the neighboring AS** check box to enable MED comparison among paths learned from confederation peers.add a new network entry. The comparison between MEDs is made only if no external autonomous systems are there in the path.
- Step 6** Check the **Treat missing MED as the least preferred one** check box to consider the missing MED attribute as having a value of infinity, making this path the least desirable; therefore, a path with a missing MED is least preferred.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.
-

Configure Policy Lists

When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. This section describes the steps required to configure policy lists.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Policy Lists**.
- Step 2** Click **Add**.
- The **Add Policy List** dialog box appears. From this dialog box, you can add a policy list name, its redistribution access (that is, permit or deny), match interfaces, specify IP addresses, match the AS path, match community names list, match metrices, and match tag numbers.
- Step 3** In the **Policy List Name** field, enter a name for the policy list.
- Step 4** Click the **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 5** Check the **Match Interfaces** check box to distribute routes that have their next hop out of one of the interfaces specified and do one of the following:
- In the **Interface** field, enter the interface name.
 - In the **Interface** field, click the ellipses to manually browse and locate the interface. Choose one or more interfaces, click **Interface**, then click **OK**.

- Step 6** In the **Specify IP** area, configure the following:
- a. Check the **Match Address** check box to redistribute any routes that have a destination network number address that is permitted by a standard access list or prefix list, and performs policy routing on packets.
Specify an access list / prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click **Access List**, then click **OK**.
 - b. Check the **Match Next Hop** check box to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified.
Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click **Access List**, then click **OK**.
 - c. Check the **Match Route Source** check box to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list.
Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click **Access List**, then click **OK**.
- Step 7** Check the **Match AS Path** check box to match a BGP autonomous system path.
Specify an AS path filter or click the ellipses to manually browse and locate the AS Path Filter. Choose one or more AS Path Filters, click **AS Path Filter**, then click **OK**.
- Step 8** Check the **Match Community Names List** check box to match a BGP community.
- a. Specify a community rule or click the ellipses to manually browse and locate the community rules. Choose one or more community rules, click **Community Rules**, then click **OK**.
 - b. Check the **Match the specified community exactly** check box to match a specific BGP community.
- Step 9** Check the **Match Metrics** check box to redistribute routes with the metric specified. If you specify more than one metric, the routes can be matched with either metric.
- Step 10** Check the **Match Tag Numbers** check box to redistribute routes in the routing table that match the specified tags. If you specify more than one tag number, routes can be matched with either metric.
- Step 11** Click **OK**.
- Step 12** Click **Apply**.

Configure AS Path Filters

An AS path filter allows you to filter the routing update message by using access lists and look at the individual prefixes within an update message. If a prefix within the update message matches the filter criteria then that individual prefix is filtered out or accepted depending on what action the filter entry has been configured to carry out. This section describes the steps required to configure AS path filters.



Note

The **as-path access-lists** are not the same as the regular firewall ACLs.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > AS Path Filters**.
- Step 2** Click **Add**.

The **Add Filter** dialog box appears. From this dialog box, you can add a filter name, its redistribution access (that is, permit or deny), and regular expression.

- Step 3** In the **Name** field, enter a name for the AS Path Filter.
 - Step 4** Click the **Permit** or **Deny** radio button to indicate the redistribution access.
 - Step 5** Specify the regular expression. Click **Build** to build regular expression.
 - Step 6** Click **Test** to test if a regular expression matches a string of your choice.
 - Step 7** Click **OK**.
 - Step 8** Click **Apply**.
-

Configure Community Rules

A community is a group of destinations that share some common attribute. You can use community lists to create groups of communities to use in a match clause of a route map. Just like an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded. This section describes the steps required to configure community rules.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Community Rules**.
- Step 2** Click **Add**.

The **Add Community Rule dialog** box appears. From this dialog box, you can add a rule name, rule type, its redistribution access (that is, permit or deny) and specific communities.
- Step 3** In the **Rule Name** field, enter a name for the community rule.
- Step 4** Click **Standard** or **Expanded** radio button to indicate the community rule type.
- Step 5** Click **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 6** To add Standard Community Rules:
 - a. In the **Communities** field, specify a community number. Valid values are from 1 to 4294967200.
 - b. (Optional) Check the **Internet** (well-known community) check box to specify the Internet community. Routes with this community are advertised to all peers (internal and external).
 - c. (Optional) Check the **Do not advertise to any peers** (well-known community) check box to specify the no-advertise community. Routes with this community are not advertised to any peer (internal or external).
 - d. (Optional) Check the **Do not export to next AS** (well-known community) check box to specify the no-export community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.
- Step 7** To add expanded community rules:
 - a. In the **Regular Expression** field, enter a regular expression. Alternately, Click **Build** to build regular expression.
 - b. Click **Test** to examine if the regular expression built, matches a string of your choice.

- Step 8** Click **OK**.
- Step 9** Click **Apply**.
-

Configure IPv4 Address Family Settings

The IPv4 settings for BGP can be set up from the IPv4 family option within the BGP configuration setup. The IPv4 family section includes subsections for General settings, Aggregate address settings, Filtering settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv4 family.

This section describes how to customize the BGP IPv4 family settings.

- [Configure IPv4 Family General Settings, page 25-9](#)
- [Configure IPv4 Family Aggregate Address Settings, page 25-10](#)
- [Configure IPv4 Family Filtering Settings, page 25-10](#)
- [Configure IPv4 Family BGP Neighbor Settings, page 25-11](#)
- [Configure IPv4 Network Settings, page 25-14](#)
- [Configure Redistribution Settings, page 25-14](#)
- [Configure Route Injection Settings, page 25-15](#)

Configure IPv4 Family General Settings

This section describes the steps required to configure the general IPv4 settings.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **General**.
- The **General IPv4 family BGP parameters** configuration pane is displayed.
- Step 3** Specify **External**, **Internal** and **Local** distances in the **Administrative Distances** area.
- Step 4** Choose a route map name from the **Learned Routes Map** drop-down list. Click **Manage** to add and configure route maps.
- Step 5** (Optional) Check the **Generate Default Route** check box to configure a BGP routing process to distribute a default route (network 0.0.0.0).
- Step 6** (Optional) Check the **Summarize subnet routes into network-level routes** check box to configure automatic summarization of subnet routes into network-level routes.
- Step 7** (Optional) Check the **Advertise inactive routes** check box to advertise routes that are not installed in the routing information base (RIB).
- Step 8** (Optional) Check the **Redistribute iBGP into an IGP** check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
- Step 9** (Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the Scanning Interval field. Valid values are from 5 to 60 seconds.

- Step 10** (Optional) Check the **Enable address tracking** check box to enable BGP next hop address tracking. Specify the delay interval between checks on updated next-hop routes installed in the routing table in the **Delay Interval** field.
- Step 11** (Optional) Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the Number of paths field and check the **iBGP multipaths** check box.
- Step 12** Click **Apply**.
-

Configure IPv4 Family Aggregate Address Settings

This section describes the steps required to define the aggregation of specific routes into one route.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Aggregate Address**.
The Aggregate Address parameters configuration pane is displayed.
- Step 3** Click **Add**.
The Add Aggregate Address pane is displayed.
- Step 4** Specify a network object in the Network field.
- Step 5** Check the **Generate autonomous system set path information** check box to generate autonomous system set path information.
- Step 6** Check the **Filters all more- specific routes from the updates** check box to filter all more-specific routes from updates.
- Step 7** Choose a route-map from the Attribute Map drop-down list. Click **Manage** to add or configure a route map.
- Step 8** Choose a route-map from the Advertise Map drop-down list. Click **Manage** to add or configure a route.
- Step 9** Choose a route-map from the Suppress Map drop-down list. Click **Manage** to add or configure a route.
- Step 10** Click **OK**.
- Step 11** Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60.
- Step 12** Click **Apply**.
-

Configure IPv4 Family Filtering Settings

This section describes the steps required to filter routes or networks received in incoming BGP updates.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.

- Step 2** Click **Filtering**.
The Define filters for BGP updates pane is displayed.
- Step 3** Click **Add**.
The Add Filter pane is displayed.
- Step 4** Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.
- Step 5** Choose an access list from the Access List drop-down list. Click **Manage** to add a new ACL.
- Step 6** Choose a protocol from the Protocol drop-down list. This is applicable only if the outbound direction is selected.
- Step 7** Choose a process ID for the protocol specified from the Process ID drop-down list.
- Step 8** Click **OK**.
- Step 9** Click **Apply**.
-

Configure IPv4 Family BGP Neighbor Settings

This section describes the steps required to define BGP neighbors and neighbor settings.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Neighbor**.
- Step 3** Click **Add**.
- Step 4** Click **General** in the left pane.
- Step 5** Enter a BGP neighbor IP address in the **IP Address** field. This IP address is added to the BGP neighbor table.
- Step 6** Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.
- Step 7** (Optional) Enter a description for the BGP neighbor in the **Description** field.
- Step 8** (Optional) Check the Shutdown neighbor administratively check box to disable a neighbor or peer group.
- Step 9** (Optional) Check the Enable address family check box to enable communication with the BGP neighbor.
- Step 10** (Optional) Check the **Global Restart Functionality for this peer** check box to enable or disable the Border Gateway Protocol (BGP) graceful restart capability for a ASA neighbor or peer group.
- Step 11** Click **Filtering** in the left pane.
- Step 12** (Optional) Choose the appropriate incoming or outgoing access control list in the Filter routes using an access list area, to distribute BGP neighbor information. Click **Manage** to add an ACL and ACEs as required.
- Step 13** (Optional) Choose the appropriate incoming or outgoing route maps in the Filter routes using a route map area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.
- Step 14** (Optional) Choose the appropriate incoming or outgoing prefix list in the Filter routes using a prefix list area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.

- Step 15** (Optional) Choose the appropriate incoming or outgoing AS path filter in the Filter routes using AS path filter area, to distribute BGP neighbor information. Click **Manage** to configure AS path filters.
- Step 16** (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box to control the number of prefixes that can be received from a neighbor.
- Enter the maximum number of prefixes allowed from a specific neighbor in the Maximum prefixes field.
 - Enter the percentage (of maximum) at which the router starts to generate a warning message in the Threshold level field. Valid values are integers between 1 to 100. The default value is 75.
 - (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:
 - Click **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field.
 - Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
- Step 17** Click **Routes** in the left pane.
- Step 18** Enter the minimum interval (in seconds) between the sending of BGP routing updates in the Advertisement Interval field.
- Step 19** (Optional) Check the **Generate Default route** check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route.
- Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the Route map drop-down list. Click **Manage** to add and configure a route map.
- Step 20** (Optional) To add conditionally advertised routes do the following:
- a. Click **Add** in the Conditionally Advertised Routes section.
 - b. Choose the route map that will be advertised from the Advertise Map drop-down list, if the conditions of the exist map or the non-exist map are met.
 - c. Do one of the following:
 - Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - d. Click **Ok**.
- Step 21** (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.
- Step 22** Click **Timers** in the left pane.
- Step 23** (Optional) Check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time and minimum hold time.
- Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. in the Keepalive frequency field. Valid values are between 0 and 65535. The default value is 60 seconds.
 - Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Hold time field. The default value is 180 seconds.

- (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Min Hold time field.

Step 24 Click **Advanced** in the left pane.

Step 25 (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.

- Choose an encryption type from the Encryption Type drop-down list.
- Enter a password in the Password field. Reenter the password in the Confirm Password field.

**Note**

The password is case-sensitive and can be up to 25 characters long, when the **service password-encryption** command is enabled and up to 81 characters long, when the **service password-encryption** command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

Step 26 (Optional) Check the **Send Community Attribute to this neighbor** check box.

Step 27 (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

Step 28 Do one of the following:

- Click **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected.
 - (Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255.
 - (Optional) Check the **Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.
- Click **Limit number of TTL hops to neighbor** to enable you to secure a BGP peering session.
 - Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254.

Step 29 (Optional) Enter a weight for the BGP neighbor connection in the Weight field.

Step 30 Choose the BGP version that the ASA will accept from the BGP version drop-down list.

**Note**

The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 31 (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.

Step 32 Choose the TCP connection mode from the TCP transport mode drop-down list.

Step 33 Click **Migration** in the left pane

Step 34 (Optional) Check the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.

- Enter the local autonomous system number in the Local AS Number field. Valid values are between 1 and 65535.

- (Optional) Check the **Do not prepend local AS number for routes received from neighbor** check box. The local AS number will not be prepended to any routes received from eBGP peer.
- (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box. The AS number from the local routing process is not prepended.
- (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box.

Step 35 Click **OK**.

Step 36 Click **Apply**.

Configure IPv4 Network Settings

This section describes the steps required to define the networks to be advertised by the BGP routing process.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Networks**.
- The Define networks to be advertised by the BGP routing process configuration pane appears.
- Step 3** Click **Add**.
- The Add Network pane is displayed.
- Step 4** Specify the network that BGP will advertise in the Address field.
- Step 5** (Optional) Choose a network or subnetmask from the Netmask drop-down list.
- Step 6** Choose a route map that should be examined to filter the networks to be advertised from the Route Map drop-down list. Click **Manage** to configure or add a route map.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.
-

Configure Redistribution Settings

This section describes the steps required to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Redistribution**.
- The Redistribution pane is displayed.
- Step 3** Click **Add**.
- The Add Redistribution pane is displayed.

- Step 4** Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.
- Step 5** Choose a process ID for the source protocol from the Process ID drop-down list.
- Step 6** (Optional) Enter a metric for the redistributed route in the Metric field.
- Step 7** Choose a route map that should be examined to filter the networks to be redistributed from the Route Map drop-down list. Click **Manage** to configure or add a route map.
- Step 8** Check one or more of the Internal, External and NSSA External Match check boxes to redistribute routes from an OSPF network.



Note This step is only applicable for redistribution from OSPF networks.

- Step 9** Click **OK**.
- Step 10** Click **Apply**.

Configure Route Injection Settings

This section describes the steps required to define the routes to be conditionally injected into the BGP routing table.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Route Injection**.
The Route Injection pane is displayed.
- Step 3** Click **Add**.
The Add Conditionally injected route pane is displayed.
- Step 4** Choose the route map that specifies the prefixes to inject into the local BGP routing table from the Inject Map drop-down list.
- Step 5** Choose the route map containing the prefixes that the BGP speaker will track from the Exist Map drop-down list.
- Step 6** Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.

Configure IPv6 Address Family Settings

The IPv6 settings for BGP can be set up from the IPv6 family option within the BGP configuration setup. The IPv6 family section includes subsections for General settings, Aggregate address settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv6 family.

This section describes how to customize the BGP IPv6 family settings.

- [Configure IPv6 Family General Settings, page 25-16](#)
- [Configure IPv6 Family Aggregate Address Settings, page 25-16](#)
- [Configure IPv6 Family BGP Neighbor Settings, page 25-17](#)
- [Configure IPv6 Network Settings, page 25-20](#)
- [Configure Redistribution Settings, page 25-14](#)
- [Configure Route Injection Settings, page 25-15](#)

Configure IPv6 Family General Settings

This section describes the steps required to configure the general IPv6 settings.

Procedure

-
Step 1	In ASDM, choose Configuration > Device Setup > Routing > BGP > IPv6 Family .
Step 2	Click General . The general IPv6 family BGP parameters configuration pane is displayed.
Step 3	Specify external, internal and local distances in the Administrative Route Distances area.
Step 4	(Optional) Check the Generate Default Route check box to configure a BGP routing process to distribute a default route (network 0.0.0.0).
Step 5	(Optional) Check the Advertise inactive routes check box to advertise routes that are not installed in the routing information base (RIB).
Step 6	(Optional) Check the Redistribute iBGP into an IGP check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
Step 7	(Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the Scanning Interval field. Valid values are from 5 to 60 seconds.
Step 8	(Optional) Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table in the Number of paths field.
Step 9	(Optional) Check the iBGP multipaths check box and specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the Number of paths field.
Step 10	Click Apply .
-

Configure IPv6 Family Aggregate Address Settings

This section describes the steps required to define the aggregation of specific routes into one route.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Aggregate Address**.
The Aggregate Address parameters configuration pane is displayed.
- Step 3** Click **Add**.
The Add Aggregate Address pane is displayed.
- Step 4** Specify an IPv6 address in the IPv6/Address Mask field. Alternately, browse to add a network object.
- Step 5** Check the **Generate autonomous system set path information** check box to generate autonomous system set path information. The path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized.
-  **Note** Do not use this form of the aggregate-address command when aggregating many paths, because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.
-
- Step 6** Check the **Filters all more-specific routes from the updates** check box to filter all more-specific routes from updates. This not only creates the aggregate route but also suppresses advertisements of more-specific routes to all neighbors.
- Step 7** Choose a route-map from the Attribute Map drop-down list. Click **Manage** to add or configure a route map. This allows attributes of the aggregate route to be changed.
- Step 8** Choose a route-map from the Advertise Map drop-down list. Click **Manage** to add or configure a route. This selects specific routes that will be used to build different components of the aggregate route.
- Step 9** Choose a route-map from the Suppress Map drop-down list. Click **Manage** to add or configure a route. This creates the aggregate route but suppresses advertisement of specified routes.
- Step 10** Click **OK**.
- Step 11** Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. This specifies the interval at which the routes will be aggregated. The default value is 30 seconds.
- Step 12** Click **Apply**.
-

Configure IPv6 Family BGP Neighbor Settings

This section describes the steps required to define BGP neighbors and neighbor settings.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Neighbor**.
- Step 3** Click **Add**.
- Step 4** Click **General** in the left pane.

- Step 5** Enter a BGP neighbor IPv6 address in the **IPv6 Address** field. This IPv6 address is added to the BGP neighbor table.
- Step 6** Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.
- Step 7** (Optional) Enter a description for the BGP neighbor in the **Description** field.
- Step 8** (Optional) Check the Shutdown neighbor administratively check box to disable a neighbor or peer group.
- Step 9** (Optional) Check the Enable address family check box to enable communication with the BGP neighbor.
- Step 10** Click **Filtering** in the left pane.
- Step 11** (Optional) Choose the appropriate incoming or outgoing route maps in the Filter routes using a route map area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.
- Step 12** (Optional) Choose the appropriate incoming or outgoing prefix list in the Filter routes using a prefix list area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.
- Step 13** (Optional) Choose the appropriate incoming or outgoing AS path filter in the Filter routes using AS path filter area, to distribute BGP neighbor information. Click **Manage** to configure AS path filters.
- Step 14** (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box to control the number of prefixes that can be received from a neighbor.
- Enter the maximum number of prefixes allowed from a specific neighbor in the Maximum prefixes field.
 - Enter the percentage (of maximum) at which the router starts to generate a warning message in the Threshold level field. Valid values are integers between 1 to 100. The default value is 75.
 - (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:
 - Click **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field.
 - Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
- Step 15** Click **Routes** in the left pane.
- Step 16** Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field.
- Step 17** (Optional) Check the **Generate Default route** check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route.
- Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the Route map drop-down list. Click **Manage** to add and configure a route map.
- Step 18** (Optional) To add conditionally advertised routes do the following:
- a. Click **Add** in the Conditionally Advertised Routes section.
 - b. Choose the route map that will be advertised from the Advertise Map drop-down list, if the conditions of the exist map or the non-exist map are met.
 - c. Do one of the following:
 - Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - d. Click **Ok**.

- Step 19** (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.
- Step 20** Click **Timers** in the left pane.
- Step 21** (Optional) Check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time and minimum hold time.
- Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. in the Keepalive frequency field. Valid values are between 0 and 65535. The default value is 60 seconds.
 - Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Hold time field. The default value is 180 seconds.
 - (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Min Hold time field.
- Step 22** Click **Advanced** in the left pane.
- Step 23** (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.
- Choose an encryption type from the Encryption Type drop-down list.
 - Enter a password in the Password field. Reenter the password in the Confirm Password field.
-  **Note** The password is case-sensitive and can be up to 25 characters long, when the **service password-encryption** command is enabled and up to 81 characters long, when the **service password-encryption** command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.
- Step 24** (Optional) Check the **Send Community Attribute to this neighbor** check box.
- Step 25** (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.
- Step 26** Do one of the following:
- Click **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected.
 - (Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255.
 - (Optional) Check the **Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.
 - Click **Limit number of TTL hops to neighbor** to enable you to secure a BGP peering session.
 - Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254.
- Step 27** (Optional) Enter a weight for the BGP neighbor connection in the Weight field.
- Step 28** Choose the BGP version that the ASA will accept from the BGP version drop-down list.



Note

The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

- Step 29** (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.
- Step 30** Choose the TCP connection mode from the TCP transport mode drop-down list.
- Step 31** Click **Migration** in the left pane
- Step 32** (Optional) Check the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.
- Enter the local autonomous system number in the Local AS Number field. Valid values are between 1 and 65535.
 - (Optional) Check the **Do not prepend local AS number for routes received from neighbor** check box. The local AS number will not be prepended to any routes received from eBGP peer.
 - (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box. The AS number from the local routing process is not prepended.
 - (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box.
- Step 33** Click **OK**.
- Step 34** Click **Apply**.

Configure IPv6 Network Settings

This section describes the steps required to define the networks to be advertised by the BGP routing process.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Networks**.
- The **Define the networks to be advertised by the BGP routing process configuration** pane appears.
- Step 3** Click **Add**.
- The **Add Network** pane is displayed.
- Step 4** In the **IPv6 Address/mask** field, specify the network that BGP will advertise.
- Step 5** Choose a route map that should be examined to filter the networks to be advertised from the **Route Map** drop-down list. Optionally, click **Manage** to configure or add a route map.
- Step 6** Click **OK**.
- Step 7** Click **Apply**.

Configure Redistribution Settings

This section describes the steps required to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Redistribution**.
- Step 3** Click **Add**.
The **Add Redistribution** pane is displayed.
- Step 4** In the **Source Protocol** drop-down list, choose the protocol from which you want to redistribute routes into the BGP domain.
- Step 5** In the **Process ID** drop-down list, choose a process ID for the source protocol. This is available only for OSPF source protocol.
- Step 6** (Optional) In the **Metric** field, enter a metric for the redistributed route.
- Step 7** In the **Route Map** drop-down list, choose a route map that should be examined to filter the networks to be redistributed. Click **Manage** to configure or add a route map.
- Step 8** Check one or more of the Match check-boxes **-Internal**, **External 1**, **External 2**, **NSSA External 1**, and **NSSA External 2** check boxes to redistribute routes from an OSPF network.



Note This step is only applicable for redistribution from OSPF networks.

- Step 9** Click **OK**.
- Step 10** Click **Apply**.
-

Configure Route Injection Settings

This section describes the steps required to define the routes to be conditionally injected into the BGP routing table.

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Route Injection**.
- Step 3** Click **Add**.
The **Add Conditionally injected route** pane is displayed.
- Step 4** In the **Inject Map** drop-down list, choose the route map that specifies the prefixes to inject into the local BGP routing table.
- Step 5** In the **Exist Map** drop-down list, choose the route map containing the prefixes that the BGP speaker will track.

- Step 6** Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.
-

Monitoring BGP

You can use the following commands to monitor the BGP 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 various BGP routing statistics, perform the following steps:



Note

To disable BGP Log messages, enter the **no bgp log-neighbor-changes** command in the router configuration mode. This disables the logging of neighbor change messages. Enter this command in router configuration mode for the BGP routing process. By default, neighbor changes are logged.

- **Monitoring > Routing > BGP Neighbors**

Each row represents one BGP neighbor. For each neighbor, the list includes the IP address, the AS number, the router ID, the state (active, idle and so on), the uptime, graceful restart capability, the restart time and the stalepath time.

- **Monitoring > Routing > BGP Routes**

Each row represents one BGP route. For each route, the list includes the status code, IP address, the next hop address, the route metric, the local preference values, the weight and the path.

History for BGP

Table 25-1 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 25-1 *Feature History for BGP*

Feature Name	Platform Releases	Feature Information
BGP Support	9.2(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Border Gateway Protocol. We introduced the following screens: Configuration > Device Setup > Routing > BGP Monitoring > Routing > BGP Neighbors, Monitoring > Routing > BGP Routes We modified the following screens: Configuration > Device Setup > Routing > Static Routes> Add > Add Static Route Configuration > Device Setup > Routing > Route Maps> Add > Add Route Map
BGP support for ASA clustering	9.3(1)	We added support for L2 and L3 clustering. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > General
BGP support for nonstop forwarding	9.3(1)	We added support for Nonstop Forwarding. We modified the following screens: Configuration > Device Setup > Routing > BGP > General, Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor, Monitoring > Routing > BGP Neighbors
BGP support for advertised maps	9.3(1)	We added support for BGPv4 advertised map. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor > Add BGP Neighbor > Routes
BGP support for IPv6	9.3(2)	We added support for IPv6. We introduced the following screen: Configuration > Device Setup > Routing > BGP > IPv6 Family



OSPF

This chapter describes how to configure the Cisco 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:

- [About OSPF, page 26-1](#)
- [Guidelines for OSPF, page 26-4](#)
- [Configure OSPFv2, page 26-6](#)
- [Configure OSPF Fast Hello Packets, page 26-7](#)
- [Customize OSPFv2, page 26-7](#)
- [Configure OSPFv3, page 26-22](#)
- [Configure Graceful Restart , page 26-33](#)
- [Example for OSPFv2, page 26-36](#)
- [Example for OSPFv3, page 26-37](#)
- [Monitoring OSPF, page 26-39](#)
- [Additional References, page 26-40](#)
- [History for OSPF, page 26-40](#)

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.

OSPF Support for Fast Hello Packets

The OSPF Support for Fast Hello Packets feature provides a way to configure the sending of hello packets in intervals less than 1 second. Such a configuration would result in faster convergence in an Open Shortest Path First (OSPF) network.

Prerequisites for OSPF Support for Fast Hello Packets

OSPF must be configured in the network already or configured at the same time as the OSPF Support for Fast Hello Packets feature.

Information About OSPF Support for Fast Hello Packets

The following sections describe concepts related to OSPF support for fast hello packets:

- [OSPF Hello Interval and Dead Interval](#)
- [OSPF Fast Hello Packets](#)
- [Benefits of OSPF Fast Hello Packets](#)

OSPF Hello Interval and Dead Interval

OSPF hello packets are packets that an OSPF process sends to its OSPF neighbors to maintain connectivity with those neighbors. The hello packets are sent at a configurable interval (in seconds). The defaults are 10 seconds for an Ethernet link and 30 seconds for a non broadcast link. Hello packets include a list of all neighbors for which a hello packet has been received within the dead interval. The dead interval is also a configurable interval (in seconds), and defaults to four times the value of the hello interval. The value of all hello intervals must be the same within a network. Likewise, the value of all dead intervals must be the same within a network.

These two intervals work together to maintain connectivity by indicating that the link is operational. If a router does not receive a hello packet from a neighbor within the dead interval, it will declare that neighbor to be down.

OSPF Fast Hello Packets

OSPF fast hello packets refer to hello packets being sent at intervals of less than 1 second. To understand fast hello packets, you should already understand the relationship between OSPF hello packets and the dead interval. See [OSPF Hello Interval and Dead Interval](#), page 26-3.

OSPF fast hello packets are achieved by using the **ospf dead-interval** command. The dead interval is set to 1 second, and the hello-multiplier value is set to the number of hello packets you want sent during that 1 second, thus providing subsecond or "fast" hello packets.

When fast hello packets are configured on the interface, the hello interval advertised in the hello packets that are sent out this interface is set to 0. The hello interval in the hello packets received over this interface is ignored.

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval.

Benefits of OSPF Fast Hello Packets

The benefit of the OSPF Fast Hello Packets feature is that your OSPF network will experience faster convergence time than it would without fast hello packets. This feature allows you to detect lost neighbors within 1 second. It is especially useful in LAN segments, where neighbor loss might not be detected by the Open System Interconnection (OSI) physical layer and data-link layer.

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.

Guidelines for OSPF

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.
- OSPFv2 supports Cisco NSF Graceful Restart and IETF NSF Graceful Restart mechanisms as defined in RFCs 4811, 4812 & 3623 respectively.
- OSPFv3 supports Graceful Restart mechanism as defined in RFC 5187.

Configure 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 [Define a Route Map, page 24-4](#). Then you generate a default route. For more information, see [Configure a Static Route, page 22-4](#).

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 [Customize OSPFv2, page 26-7](#).

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:

Procedure

-
- 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 Spanned EtherChannel or Individual Interface 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 [Configure OSPFv2 Area Parameters, page 26-14](#) 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 [Configure Route Summarization Between OSPFv2 Areas, page 26-11](#) for more information.

Configure OSPF Fast Hello Packets

This section describes how to configure OSPF Fast Hello Packets.

Procedure

Customize OSPFv2

This section explains how to customize the OSPFv2 processes.

- [Redistribute Routes Into OSPFv2, page 26-7](#)
- [Configure Route Summarization When Redistributing Routes Into OSPFv2, page 26-9](#)
- [Configure Route Summarization Between OSPFv2 Areas, page 26-11](#)
- [Configure OSPFv2 Interface Parameters, page 26-12](#)
- [Configure OSPFv2 Area Parameters, page 26-14](#)
- [Configure an OSPFv2 NSSA, page 26-15](#)
- [Configure an IP Address Pool for Clustering \(OSPFv2 and OSPFv3\), page 26-16](#)
- [Define Static OSPFv2 Neighbors, page 26-18](#)
- [Configure Route Calculation Timers, page 26-19](#)
- [Log Neighbors Going Up or Down, page 26-19](#)
- [Configure Filtering in OSPF, page 26-20](#)
- [Configure a Virtual Link in OSPF, page 26-21](#)

Redistribute 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 [Configure a Static Route, page 22-4](#), and then define a route map according to [Define a Route Map, page 24-4](#).

To redistribute static, connected, RIP, or OSPFv2 routes into an OSPFv2 process, perform the following steps:

Procedure

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.
- **BGP**—Redistribute routes from the BGP 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 [Define a Route Map, page 24-4](#).
- Step 13** Click **OK**.
-

Configure 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.

To configure route summarization, you can:

- [Add a Route Summary Address, page 26-10](#)
- [Add or Edit an OSPF Summary Address, page 26-10](#)

Add 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:

Procedure

-
- 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**.
-

Add or Edit an OSPF Summary Address

Procedure

-
- 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.
-

Configure 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:

Procedure

-
- 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.
-

Configure OSPFv2 Interface Parameters

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the Hello interval, the Dead interval, and the Authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

To configure OSPFv2 interface parameters, perform the following steps:

Procedure

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 [Configure OSPFv2 Area Parameters, page 26-14](#) 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 [Define Static OSPFv2 Neighbors, page 26-18](#) for more information.
- You need to define a static route pointing to the crypto endpoint. See [Configure a Static Route, page 22-4](#) 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 In the Intervals section, 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 8192 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 8192 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 8192 seconds. The default value is 1 second.

Step 15 In the Detecting Lost Neighbors section, do one of the following:

- Click **Configure interval within which hello packets are not received before the router declares the neighbor to be down**. In the Dead Interval field, 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 8192 seconds. The default value of this setting is four times the interval that was set in the Hello Interval field.
- Click **Send fast hello packets within 1 seconds dead interval**. In the Hello multiplier field, specify the number of hello packets to be sent per second. Valid values are between 3 and 20.

Configure 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.

Procedure

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**.
-

Configure 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.

Procedure

-
- 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**.
-

Configure 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 Individual Interface clustering.

Procedure

To assign a range of IPv4 addresses for the router ID cluster pool in Individual Interface 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 Individual Interface 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**.

Define 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 22, “Static and Default Routes,”](#) for more information about creating static routes.

Procedure

- 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**.
-

Configure 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.

Procedure

-
- 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 Initial SPF Delay, specifies the time (in milliseconds) between when OSPF receives a topology change and when the SPF calculation starts. Valid values range from 0 to 600000 milliseconds.
 - The Minimum SPF Hold Time, specifies the hold time (in milliseconds) between consecutive SPF calculations. Valid values range from 0 to 600000 milliseconds.
 - The Maximum SPF Wait Time, specifies the maximum wait time between two consecutive SPF calculations. Valid values range from 0 to 600000 milliseconds.
- Step 5** Click **OK**.
-

Log Neighbors Going Up or Down

By default, a syslog message is generated when an OSPFv2 neighbor goes up or down.

Procedure

-
- 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.

Configure Filtering in OSPF

The Filtering pane displays the ABR Type 3 LSA filters that have been configured for each OSPF process.

ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.



Note Only Type 3 LSAs that originate from an ABR are filtered.

To configure filtering in OSPF, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Filtering**.
- Step 2** Click **Add** or **Edit**.
The Add or Edit OSPF Filtering Entry dialog box lets you add new filters to the Filter table or modify an existing filter. Some of the filtering information cannot be changed when you edit an existing filter.
- Step 3** Choose the OSPF process that is associated with the filter entry from the OSPF Process drop-down list.
- Step 4** Choose the Area ID that is associated with the filter entry from the Area ID drop-down list. If you are editing an existing filter entry, you cannot modify this setting.
- Step 5** Choose a prefix list from the Prefix List drop-down list.
- Step 6** Choose the traffic direction being filtered from the Traffic Direction drop-down list.
Choose Inbound to filter LSAs coming into an OSPF area, or Outbound to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.

- Step 7** Click **Manage** to display the Configure Prefix Lists dialog box, from which you can add, edit, or delete prefix lists and prefix rules. For more information, see [Configure Prefix Lists, page 24-8](#) and the [Configure the Metric Values for a Route Action, page 24-9](#).
- Step 8** Click **OK**.
-

Configure 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 [Configure OSPFv2 Area Parameters, page 26-14](#) 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:
- 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.
 - In the Key field, enter an alphanumeric character string of up to 16 bytes. The Key ID displays for the selected interface.
 - 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**.
-

Configure OSPFv3

This section describes how to configure OSPFv3 routing processes.

- [Enable OSPFv3, page 26-23](#)
- [Configure OSPFv3 Interface Parameters, page 26-23](#)
- [Configure OSPFv3 Area Parameters, page 26-25](#)
- [Configure a Virtual Link Neighbor, page 26-26](#)
- [Configure OSPFv3 Passive Interfaces, page 26-27](#)
- [Configure OSPFv3 Administrative Distance, page 26-27](#)
- [Configure OSPFv3 Timers, page 26-28](#)
- [Define Static OSPFv3 Neighbors, page 26-29](#)
- [Send Syslog Messages, page 26-30](#)

- [Suppress Syslog Messages, page 26-30](#)
- [Calculate Summary Route Costs, page 26-30](#)
- [Generate a Default External Route into an OSPFv3 Routing Domain, page 26-31](#)
- [Configure an IPv6 Summary Prefix, page 26-31](#)
- [Redistribute IPv6 Routes, page 26-32](#)

Enable 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.

Procedure

-
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 Configure OSPFv3 Area Parameters, page 26-25 .
-

Configure OSPFv3 Interface Parameters

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

Procedure

-
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.
-

Configure OSPFv3 Area Parameters

Procedure

-
- 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 [Configure a Virtual Link Neighbor, page 26-26](#).
-

Configure 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.

Configure OSPFv3 Passive Interfaces

Procedure

-
- 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.
-

Configure OSPFv3 Administrative Distance

Procedure

-
- 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.

Configure OSPFv3 Timers

You can set LSA arrival, LSA pacing, and throttling timers for OSPFv3.

Procedure

-
- 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.

Define 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 22, “Static and Default Routes,”](#) for more information about creating static routes.

Procedure

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**.

Send Syslog Messages

Configure the router to send a syslog message when an OSPFv3 neighbor goes up or down.

Procedure

-
- 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.
-

Suppress 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:

Procedure

-
- 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**.
-

Calculate Summary Route Costs

Procedure

-
- 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**.
-

Generate a Default External Route into an OSPFv3 Routing Domain

Procedure

-
- 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.
 - Check the **Always advertise** check box to always advertise the default route, whether or not one exists.
 - 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.
 - 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 routeThe default is the Type 2 external route.
 - 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.
-

Configure an IPv6 Summary Prefix

Procedure

-
- 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.
-

Redistribute IPv6 Routes

Procedure

-
- 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.

Configure Graceful Restart

The ASA may experience some known failure situations, that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a component failure (i.e., active unit crash with standby unit taking over in failover (HA) mode, master unit crash with slave unit elected as new master in cluster mode), or when there is a scheduled hitless software upgrade.

Graceful restart is supported on both OSPFv2 and OSPFv3. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623). You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.



Note

When fast hellos are configured for OSPFv2, graceful restart does not occur when the active unit reloads and the standby unit becomes active. This is because the time taken for the role change is more than the configured dead interval.

Configuring Graceful Restart for OSPFv2

There are two graceful restart mechanisms for OSPFv2, Cisco NSF and IETF NSF. Only one of these graceful restart mechanisms can be configured at a time for an ospf instance. An NSF-aware device can be configured as both Cisco NSF helper and IETF NSF helper but a NSF-capable device can be configured in either Cisco NSF or IETF NSF mode at a time for an ospf instance.

Configure Cisco NSF Graceful Restart for OSPFv2

Configure Cisco NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties**.
 - Step 2** Under Configuring Cisco NSF, check the **Enable Cisco nonstop forwarding (NSF)** check box.
 - Step 3** (Optional) Check the **Cancels NSF restart when non-NSF-aware neighboring networking devices are detected** check box if required.
 - Step 4** (Optional) Under Configuring Cisco NSF helper, uncheck the **Enable Cisco nonstop forwarding (NSF) for helper mode** check box.
-
-  **Note** This is checked by default. Uncheck this to disable the Cisco NSF helper mode on NSF-aware device.
-
- Step 5** Click **OK**.
 - Step 6** Click **Apply** to save your changes.
-

Configure IETF NSF Graceful Restart for OSPFv2

Configure IETF NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties**.
 - Step 2** Under Configuring IETF NSF, check the **Enable IETF nonstop forwarding (NSF)** check box.
 - Step 3** (Optional) Enter the restart interval in seconds in the Length of graceful restart interval field.
-
-  **Note** The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
-
- Step 4** (Optional) Under Configuring IETF NSF helper, uncheck the **Enable IETF nonstop forwarding (NSF) for helper mode** check box.

**Note**

This is checked by default. Uncheck this to disable the IETF NSF helper mode on NSF-aware device.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configuring Graceful Restart for OSPFv3

Configuring the NSF graceful-restart feature for OSPFv3 involves two steps; configuring a device to be NSF-capable and then configuring a device to be NSF-aware.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup > Advanced > Add NSF Properties**.

Step 2 Under Configuring Graceful Restart, check the **Enable Graceful Restart** check box.

Step 3 (Optional) Enter a value for the restart interval in the Restart Interval field.

**Note**

The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 4 Under Configuring Graceful Restart Helper, check the **Enable Graceful Restart Helper** check box. .

**Note**

This is checked by default. Uncheck this to disable the Graceful-restart helper mode on a NSF-aware device.

Step 5 (Optional) Check the **Enable LSA checking** check box to enable strict link state advertisement checking.

**Note**

When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.

Step 6 Click **OK**.

Step 7 Click **Apply** to save your changes.

Remove the OSPF Configuration

Remove the OSPFv2 configuration.

Procedure

-
- 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**.
-

Remove the OSPFv3 configuration.

Procedure

-
- 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**.
-

Example for OSPFv2

The following example shows how to enable and configure OSPFv2 with various optional processes:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.
- Step 2** Click the **Process Instances** tab and in the OSPF Process 1 field, type **2**.
- Step 3** Click the **Area/Networks** tab, and click **Add**.
- Step 4** Enter **0** in the Area ID field.
- Step 5** In the Area Networks area, enter **10.0.0.0** in the IP Address field.
- Step 6** Choose 255.0.0.0 from the Netmask drop-down list.
- Step 7** Click **OK**.
- Step 8** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Redistribution**.
- Step 9** Click **Add**.
- The Add/Edit OSPF Redistribution Entry dialog box appears.
- Step 10** In the Protocol area, click the **OSPF** radio button to choose the source protocol from which the routes are being redistributed. Choosing OSPF redistributes routes from another OSPF routing process.
- Step 11** Choose the OSPF process ID from the OSPF Process drop-down list.
- Step 12** In the Match area, check the **Internal** check box.
- Step 13** In the Metric Value field, enter **5** for the metric value for the routes being redistributed.
- Step 14** From the Metric Type drop-down list, choose 1 for the Metric Type value.
- Step 15** From the Route Map drop-down list, choose 1.
- Step 16** Click **OK**.

- Step 17** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Interface**.
- Step 18** From the Properties tab, choose the **inside** interface and click **Edit**.
The Edit OSPF Properties dialog box appears.
- Step 19** In the Cost field, enter **20**.
- Step 20** Click **Advanced**.
- Step 21** In the Retransmit Interval field, enter **15**.
- Step 22** In the Transmit Delay field, enter **20**.
- Step 23** In the Hello Interval field, enter **10**.
- Step 24** In the Dead Interval field, enter **40**.
- Step 25** Click **OK**.
- Step 26** In the Edit OSPF Properties dialog box, enter **20** in the Priorities field, and click **OK**.
- Step 27** Click the **Authentication** tab.
The Edit OSPF Authentication dialog box appears.
- Step 28** In the Authentication area, click the **MD5** radio button.
- Step 29** In the MD5 and Key ID area, enter **cisco** in the MD5 Key field, and **1** in the MD5 Key ID field.
- Step 30** Click **OK**.
- Step 31** Choose **Configuration > Device Setup > Routing > OSPF > Setup**, and click the **Area/Networks** tab.
- Step 32** Choose the **OSPF 2** process and click **Edit**.
The Edit OSPF Area dialog box appears.
- Step 33** In the Area Type area, choose **Stub**.
- Step 34** In the Authentication area, choose **None**, and enter **20** in the Default Cost field.
- Step 35** Click **OK**.
- Step 36** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup**.
- Step 37** Click the **Process Instances** tab and check the **OSPF process 2** check box.
- Step 38** Click **Advanced**.
The Edit OSPF Area dialog box appears.
- Step 39** In the Timers area, enter **10** in the SPF Delay Time field and **20** in the SPF Hold Time field.
- Step 40** In the Adjacency Changes area, check the **Log Adjacency Change Details** check box.
- Step 41** Click **OK**.
- Step 42** Click **Reset**.

Example for OSPFv3

The following example shows how to configure OSPFv3 routing in ASDM:

- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup**.

- Step 2** On the Process Instances tab, do the following:
- a. Check the **Enable OSPFv3 Process** check box.
 - b. Enter **1** in the Process ID field.
- Step 3** Click the **Areas** tab, then click **Add** to display the Add OSPFv3 Area dialog box.
- Step 4** From the OSPFv3 Process ID drop-down list, choose **1**.
- Step 5** Enter **22** in the Area ID field.
- Step 6** Choose **Normal** from the Area Type drop-down list.
- Step 7** Enter **10** in the Default Cost field.
- Step 8** Check the **Redistribution imports routes to normal and NSSA areas** check box.
- Step 9** Enter **20** in the Metric field.
- Step 10** Choose **1** from the Metric Type drop-down list.
- Step 11** Check the **inside** check box as the specified interface being used.
- Step 12** Check the **Enable Authentication** check box.
- Step 13** Enter **300** in the Security Policy Index field.
- Step 14** Choose **SHA-1** from the Authentication Algorithm drop-down list.
- Step 15** Enter **12345ABCDE** in the Authentication Key field.
- Step 16** Choose **DES** from the Encryption Algorithm drop-down list.
- Step 17** Enter **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.
-

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can also use the information provided to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your device packets are taking through the network.

To monitor or display various OSPFv2 routing statistics in ASDM, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Monitoring > Routing > OSPF LSAs**.
- Step 2** You can select and monitor OSPF LSAs, Types 1 through 5 and 7. Each pane shows one LSA type, as follows:
- Type 1 LSAs represent the routes in an area under a process.
 - Type 2 LSAs show the IP address of the designated router that advertises the routers.
 - Type 3 LSAs show the IP address of the destination network.
 - Type 4 LSAs show the IP address of the AS boundary router.
 - Type 5 LSAs and Type 7 LSAs show the IP address of the AS external network.
- Step 3** Click **Refresh** to update each LSA type pane.
- Step 4** In the main ASDM window, choose **Monitoring > Routing > OSPF Neighbors**.
- In the OSPF Neighbors pane, each row represents one OSPF neighbor. In addition, the OSPF Neighbors pane shows the network on which the neighbor is running, the priority, the state, the amount of dead time in seconds, the IP address of the neighbor, and the interface on which it is running. For a list of possible states for an OSPF neighbor, see RFC 2328.
- Step 5** Click **Refresh** to update the OSPF Neighbors pane.
-

To monitor or display various OSPFv3 routing statistics in ASDM, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Monitoring > Routing > OSPFv3 LSAs**.
- Step 2** You can select and monitor OSPFv3 LSAs. Choose a link-state type to display its status according to specified parameters from the Link State type drop-down list. The supported link-state types are router, network, inter-area prefix, inter-area router, AS external, NSSA, link, and intra-area prefix.
- Step 3** Click **Refresh** to update each link-state type.

Step 4 In the main ASDM window, choose **Monitoring > Routing > OSPFv3 Neighbors**.

In the OSPFv3 Neighbors pane, each row represents one OSPFv3 neighbor. In addition, the OSPFv3 Neighbors pane shows the IP address of the neighbor, the priority, the state, the amount of dead time in seconds, and the interface on which it is running. For a list of possible states for an OSPFv3 neighbor, see RFC 5340.

Step 5 Click **Refresh** to update the OSPFv3 Neighbors pane.

Additional References

RFCs

RFC	Title
2328	OSPFv2
4552	OSPFv3 Authentication
5340	OSPF for IPv6

History for OSPF

Table 26-1 Feature History for OSPF

Feature Name	Platform Releases	Feature Information
OSPF Support	7.0(1)	Support was added for route data, authentication, and redistribution and monitoring of routing information using the Open Shortest Path First (OSPF) routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > OSPF.

Table 26-1 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 Spanned EtherChannel load balancing are supported in the clustering environment.
OSPFv3 Support for IPv6		OSPFv3 routing is supported for IPv6. We introduced the following screens: Configuration > Device Setup > Routing > OSPFv3 > Setup, Configuration > Device Setup > Routing > OSPFv3 > Interface, Configuration > Device Setup > Routing > OSPFv3 > Redistribution, Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix, Configuration > Device Setup > Routing > OSPFv3 > Virtual Link, Monitoring > Routing > OSPFv3 LSAs, Monitoring > Routing > OSPFv3 Neighbors.
OSPF support for Fast Hellos	9.2(1)	OSPF supports the Fast Hello Packets feature, resulting in a configuration that results in faster convergence in an OSPF network. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Interface > Edit OSPF Interface Advanced Properties
Timers		New OSPF timers were added; old ones were deprecated. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup > Edit OSPF Process Advanced Properties
Route filtering using access-list		Route filtering using ACL is now supported. We introduced the following screen: Configuration > Device Setup > Routing > OSPF > Filtering Rules > Add Filter Rules
OSPF Monitoring enhancements		Additional OSPF monitoring information was added.
OSPF redistribute BGP		OSPF redistribution feature was added. We added the following screen: Configuration > Device Setup > Routing > OSPF > Redistribution
OSPF Support for Non-Stop Forwarding (NSF)	9.3(1)	OSPFv2 and OSPFv3 support for NSF was added. We added the following screens: Configuration > Device Setup > Routing > OSPF > Setup > NSF Properties, Configuration > Device Setup > Routing > OSPFv3 > Setup > NSF Properties



EIGRP

This chapter describes how to configure the Cisco ASA to route data, perform authentication, and redistribute routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP).

- [About EIGRP, page 27-1](#)
- [Guidelines for EIGRP, page 27-2](#)
- [Configure an EIGRP Process, page 27-3](#)
- [Configure EIGRP, page 27-3](#)
- [Customize EIGRP, page 27-6](#)
- [Monitoring for EIGRP, page 27-18](#)
- [History for EIGRP, page 27-19](#)

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.

Guidelines for EIGRP

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 Spanned EtherChannel and Individual Interface clustering when configured to use both EIGRP and OSPFv2.
- In a Individual Interface 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.

Configure an EIGRP Process

To configure EIGRP routing on the ASA, perform the following steps:

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > EIGRP .
Step 2	Enable the EIGRP routing process by checking the Enable this EIGRP process check box on the Process Instances tab. See Enable EIGRP, page 27-4 or the Enable EIGRP Stub Routing, page 27-5 .
Step 3	Define the networks and interfaces that will participate in EIGRP routing on the Setup > Networks tab. See Define a Network for an EIGRP Routing Process, page 27-6 for more information.
Step 4	(Optional) Define route filters on the Filter Rules pane. Route filtering provides more control over the routes that are allowed to be sent or received in EIGRP updates. See Filter Networks in EIGRP, page 27-13 for more information.
Step 5	(Optional) Define route redistribution in the Redistribution pane. You can redistribute routes discovered by RIP and OSPF to the EIGRP routing process. You can also redistribute static and connected routes to the EIGRP routing process. See Redistribute Routes Into EIGRP, page 27-12 for more information.
Step 6	(Optional) Define static EIGRP neighbors on the Static Neighbor pane. See Define an EIGRP Neighbor, page 27-11 for more information.
Step 7	(Optional) Define summary addresses on the Summary Address pane. See Configure the Summary Aggregate Addresses on Interfaces, page 27-9 for more information about defining summary addresses.
Step 8	(Optional) Define interface-specific EIGRP parameters on the Interfaces pane. These parameters include EIGRP message authentication, hold time, hello interval, delay metric, and the use of split-horizon. See Configure Interfaces for EIGRP, page 27-7 for more information.
Step 9	(Optional) Control the sending and receiving of default route information in EIGRP updates on the Default Information pane. By default, default routes are sent and accepted. See Configure Default Information in EIGRP, page 27-16 for more information.
-

Configure 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.

- [Enable EIGRP, page 27-4](#)
- [Enable EIGRP Stub Routing, page 27-5](#)

Enable EIGRP

You can only enable one EIGRP routing process on the ASA.

To enable EIGRP, perform the following steps:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.
The EIGRP Setup pane appears.
The three tabs on the main EIGRP Setup pane used to enable EIGRP are as follows:
- The Process Instances tablets you enable an EIGRP routing process for each context. Single context mode and multiple context mode are both supported. See [Enable EIGRP, page 27-4](#) and the [Enable EIGRP Stub Routing, page 27-5](#) for more information.
 - The Networks tablets you specify the networks used by the EIGRP routing process. For an interface to participate in EIGRP routing, it must fall within the range of addresses defined by the network entries. For directly connected and static networks to be advertised, they must also fall within the range of the network entries. See [Define a Network for an EIGRP Routing Process, page 27-6](#) for more information.
 - The Passive Interfaces tablets you configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates. The Passive Interface table lists each interface that is configured as a passive interface.
- Step 2** Check the **Enable this EIGRP process** check box.
You can only enable one EIGRP routing process on the device. You must enter an autonomous system number (AS) for the routing process in the EIGRP Process field before you can save your changes.
- Step 3** In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
- Step 4** (Optional) Click **Advanced** to configure the EIGRP process settings, such as the router ID, default metrics, stub routing, neighbor changes, and the administrative distances for the EIGRP routes.
- Step 5** Click the **Networks** tab.
- Step 6** To add a new network entry, click **Add**.
The Add EIGRP Network dialog box appears. To remove a network entry, choose an entry in the table and click **Delete**.
- Step 7** Choose the AS number of the EIGRP routing process from the drop-down list.
- Step 8** Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.
-  **Note** To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.
-
- Step 9** Enter a network mask to apply to the IP address in the Network Mask field.
- Step 10** Click **OK**.
-

Enable 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:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > EIGRP > Setup . The EIGRP Setup pane appears.
Step 2	Check the Enable EIGRP routing check box.
Step 3	In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
Step 4	Click Advanced to configure the EIGRP stub routing process. The Edit EIGRP Process Advanced Properties dialog box appears.
Step 5	In the Stub area on the Edit EIGRP Process Advanced Properties dialog box, choose one or more of the following EIGRP stub routing processes: • Stub Receive only—Configures the EIGRP stub routing process to receive route information from the neighbor routers but does not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options. • Stub Connected—Advertises connected routes. • Stub Static—Advertises static routes. • Stub Redistributed—Advertises redistributed routes. • Stub Summary—Advertises summary routes.
Step 6	Click OK .
Step 7	Click the Networks tab.
Step 8	Click Add to add a new network entry. The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click Delete .
Step 9	Choose the AS number of the EIGRP routing process from the drop-down list.
Step 10	Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.

**Note**

To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 11 Enter a network mask to apply to the IP address in the Network Mask field.

Step 12 Click **OK**.

Customize EIGRP

This section describes how to customize the EIGRP routing.

- [Define a Network for an EIGRP Routing Process, page 27-6](#)
- [Configure Interfaces for EIGRP, page 27-7](#)
- [Configure Passive Interfaces, page 27-8](#)
- [Configure the Summary Aggregate Addresses on Interfaces, page 27-9](#)
- [Change the Interface Delay Value, page 27-10](#)
- [Enable EIGRP Authentication on an Interface, page 27-10](#)
- [Define an EIGRP Neighbor, page 27-11](#)
- [Redistribute Routes Into EIGRP, page 27-12](#)
- [Filter Networks in EIGRP, page 27-13](#)
- [Customize the EIGRP Hello Interval and Hold Time, page 27-15](#)
- [Disable Automatic Route Summarization, page 27-15](#)
- [Configure Default Information in EIGRP, page 27-16](#)
- [Disable EIGRP Split Horizon, page 27-17](#)
- [Restart the EIGRP Process, page 27-18](#)

Define 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:

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
- Step 4** Click the **Networks** tab.
- Step 5** Click **Add** to add a new network entry.
- The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.
- Step 6** Choose the AS number of the EIGRP routing process from the drop-down list.
- Step 7** Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.



Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

- Step 8** Enter a network mask to apply to the IP address in the Network Mask field.
- Step 9** Click **OK**.

Configure Interfaces for EIGRP

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, you can configure the ASA that includes the network to which the interface is attached, and prevent that interface from sending or receiving EIGRP updates.

To configure interfaces for EIGRP, perform the following steps:

Procedure

- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**. The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** Click **OK**.
- Step 4** Choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.
- The Interface pane appears and displays the EIGRP interface configurations. The Interface Parameters table displays all of the interfaces on the ASA and lets you modify the following settings on a per-interface basis:
- Authentication key and mode.
 - The EIGRP hello interval and hold time.
 - The interface delay metric used in EIGRP metric calculations.
 - The use of split-horizon on the interface.
- Step 5** Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The Edit EIGRP Interface Entry dialog box appears.

- Step 6** In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.
- Step 7** In the Hello Interval field, enter the interval between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
- Step 8** In the Hold Time field, enter the hold time, in seconds. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
- Step 9** Check the **Enable** check box for Split Horizon.
- Step 10** In the Delay field, enter the delay value. The delay time is in tens of microseconds. Valid values range from 1 to 16777215.
- Step 11** Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages.
- Step 12** Enter the Key or Key ID values.
- In the Key field, enter the key to authenticate EIGRP updates. The key can contain up to 16 characters.
 - In the Key ID field, enter the key identification value. Valid values range from 1 to 255.
- Step 13** Click **OK**.
-

Configure Passive Interfaces

You can configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates. In ASDM, the Passive Interface table lists each interface that is configured as a passive interface.

To configure passive interfaces, perform the following steps:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**. The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** Click **OK**.
- Step 4** Click the **Passive Interfaces** tab.
- Step 5** Choose the interface that you want to configure from the drop-down list.
- Step 6** Check the **Suppress routing updates on all interfaces** check box to specify all interfaces as passive. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when the check box is checked.
- Step 7** Click **Add** to add a passive interface entry.
- The Add EIGRP Passive Interface dialog box appears. Choose the interface that you want to make passive and click **Add**. To remove a passive interface, choose the interface in the table and click **Delete**.

Step 8 Click **OK**.

Configure 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**. The Interface pane shows the EIGRP interface configurations. The Interface Parameters table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see [Configure Interfaces for EIGRP, page 27-7](#).
- Step 2** To configure the EIGRP parameters for an interface, double-click an interface entry or select the entry and click **Edit**.
- Step 3** Click **OK**.
- Step 4** Choose **Configuration > Device Setup > Routing > EIGRP > Summary Address**. The Summary Address pane displays a table of the statically-defined EIGRP summary addresses. By default, EIGRP summarizes subnet routes to the network level. You can create statically defined EIGRP summary addresses to the subnet level from the Summary Address pane.
- Step 5** Click **Add** to add a new EIGRP summary address, or to click **Edit** to edit an existing EIGRP summary address in the table. The Add Summary Address or Edit Summary Address dialog box appears. You can also double-click an entry in the table to edit that entry.
- Step 6** In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
- Step 7** In the Interface drop-down list, choose the interface from which the summary address is advertised.
- Step 8** In the IP Address field, enter the IP address of the summary route.
- Step 9** In the Netmask field, choose or enter the network mask to apply to the IP address.
- Step 10** Enter the administrative distance for the route in the Administrative Distance field. If left blank, the route has the default administrative distance of 5.
- Step 11** Click **OK**.
-

Change 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**. The Interface pane shows the EIGRP interface configurations. The Interface Parameters table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see [Configure Interfaces for EIGRP, page 27-7](#).
- Step 2** Double-click an interface entry or choose the Interface entry and click **Edit** to configure the delay value in the EIGRP parameters for an interface. The Edit EIGRP Interface Entry dialog box appears.
- Step 3** In the Delay field, enter the delay time, which is in tens of microseconds. Valid values are from 1 to 16777215.
- Step 4** Click **OK**.
-

Enable 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**. The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can range from 1 to 65535.
- Step 4** Click the **Networks** tab.
- Step 5** Click **Add** to add a new network entry.

The Add EIGRP Network dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 6 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 7 In the IP Address field, enter the IP address of the networks to participate in the EIGRP routing process.



Note To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 8 In the Network Mask field, choose or enter a network mask to apply to the IP address.

Step 9 Click **OK**.

Step 10 Choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**.

The Interface pane displays the EIGRP interface configurations. The Interface Parameters table displays all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see [Configure Interfaces for EIGRP, page 27-7](#).

Step 11 Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages. After you check this check box, provide one of the following:

- In the Key field, enter the key to authenticate EIGRP updates. The key can include up to 16 characters.
- In the Key ID field, enter the key identification value. Valid values range from 1 to 255.

Step 12 Click **OK**.

Define 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:

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.

The EIGRP Setup pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration > Device Setup > Routing > EIGRP > Static Neighbor**.

The Static Neighbor pane appears and displays the statically-defined EIGRP neighbors. An EIGRP neighbor sends EIGRP routing information to and receives EIGRP routing information from the ASA. Normally, neighbors are dynamically discovered through the neighbor discovery process. However, on point-to-point, nonbroadcast networks, you must statically define the neighbors.

Each row of the Static Neighbor table displays the EIGRP autonomous system number for the neighbor, the neighbor IP address, and the interface through which the neighbor is available.

From the Static Neighbor pane, you can add or edit a static neighbor.

- Step 5** Click **Add** or **Edit** to add or edit a EIGRP static neighbor.
- The Add or Edit EIGRP Neighbor Entry dialog box appears.
- Step 6** Choose the EIGRP AS number from the drop-down list for the EIGRP process for which the neighbor is being configured.
- Step 7** Choose the Interface Name from the Interface Name drop-down list, which is the interface through which the neighbor is available.
- Step 8** Enter the IP address of the neighbor in the Neighbor IP Address field.
- Step 9** Click **OK**.
-

Redistribute 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 24, “Route Maps,”](#) for more information about creating a route map.

To redistribute routes into the EIGRP routing process, perform the following steps:

Procedure

- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.
- The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.
- Step 4** Choose **Configuration > Device Setup > Routing > EIGRP > Redistribution**.
- The Redistribution pane displays the rules for redistributing routes from other routing protocols to the EIGRP routing process. When redistributing static and connected routes to the EIGRP routing process, metrics are not required to be configured, although this is recommended. Each row of the Redistribution pane table includes a route redistribution entry.
- Step 5** Click **Add** to add a new redistribution rule. If you are editing an existing redistribution rule, go to Step 6.
- The Add EIGRP Redistribution Entry dialog box appears.
- Step 6** Choose the address in the table and click **Edit** to edit an existing EIGRP static neighbor. You can also double-click an entry in the table to edit that entry.
- The Edit EIGRP Redistribution Entry dialog box appears.

- Step 7** Choose the AS number of the EIGRP routing process to which the entry applies from the drop-down list.
- Step 8** In the Protocol area, click the radio button next to one of the following protocols for the routing process:
- **Static** to redistribute static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
 - **Connected** to redistribute connected routes into the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
 - **RIP** to redistributes routes discovered by the RIP routing process to EIGRP.
 - **OSPF** to redistribute routes discovered by the OSPF routing process to EIGRP.
- Step 9** In the Optional Metrics area, choose one of the following metrics used for the redistributed route:
- **Bandwidth**, which is the EIGRP bandwidth metric in kilobits per second. Valid values range from 1 to 4294967295.
 - **Delay**, which is the EIGRP delay metric, in 10-microsecond units. Valid values range from 0 to 4294967295.
 - **Reliability**, which is the EIGRP reliability metric. Valid values range from 0 to 255; 255 indicates 100 percent reliability.
 - **Loading**, which is the EIGRP effective bandwidth (loading) metric. Valid values range from 1 to 255; 255 indicates 100 percent loaded.
 - **MTU**, which is the MTU of the path. Valid values range from 1 to 65535.
- Step 10** Choose the route map from the Route Map drop-down list to define which routes are redistributed into the EIGRP routing process. For more details about how to configure a route map, see [Chapter 24, “Route Maps.”](#)
- Step 11** In the Optional OSPF Redistribution area, click one of the following OSPF radio buttons to further specify which OSPF routes are redistributed into the EIGRP routing process:
- **Match Internal** to match routes internal to the specified OSPF process.
 - **Match External 1** to match type 1 routes external to the specified OSPF process.
 - **Match External 2** to match type 2 routes external to the specified OSPF process.
 - **Match NSSA-External 1** to match type 1 routes external to the specified OSPF NSSA.
 - **Match NSSA-External 2** to match type 2 routes external to the specified OSPF NSSA.
- Step 12** Click **OK**.
-

Filter 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.
The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** In the EIGRP Process field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.
- Step 4** Choose **Configuration > Device Setup > Routing > EIGRP > Filter Rules**.
The Filter Rules pane appears and displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRP routing process.
Each row of the Filter Rule table describes a filter rule for a specific interface or routing protocol. For example, a filter rule with a direction of in on the outside interface would apply filtering to any EIGRP updates received on the outside interface. A filter rule with a direction of out with OSPF 10 specified as the routing protocol would apply the filter rules to routes redistributed into the EIGRP routing process in outbound EIGRP updates.
- Step 5** Click **Add** to add a filter rule. If you are editing an already existing filter rule, skip to Step 6.
The Add Filter Rules dialog box appears.
- Step 6** To edit a filter rule, choose the filter rule in the table and click **Edit**.
The Edit Filter Rules dialog appears. You can also double-click a filter rule to edit the rule. To remove a filter rule, choose the filter rule in the table and click **Delete**.
- Step 7** Choose the AS number from the drop-down list of the EIGRP routing process to which the entry applies.
- Step 8** Choose the direction of the filter routes from the drop-down list.
Choose **in** for rules that filter routes from incoming EIGRP routing updates. Choose **out** to filter routes from EIGRP routing updates that are sent by the ASA.
If you choose **out**, the Routing process field becomes active. Choose the type of route to be filtered. You can filter routes redistributed from static, connected, RIP, and OSPF routing processes. Filters that specify a routing process filter those routes from updates sent on all interfaces.
- Step 9** Enter the OSPF process ID in the ID field.
- Step 10** Click the **Interface** radio button and choose the interface to which the filter applies.
- Step 11** Click **Add** or **Edit** to define an ACL for the filter rule. Clicking **Edit** opens the Network Rule dialog box for the selected network rule.
The Network Rule dialog box appears.
- Step 12** In the Action drop-down list, choose Permit to allow the specified network to be advertised; choose Deny to prevent the specified network from being advertised.
- Step 13** In the IP Address field, type IP address of the network being permitted or denied. To permit or deny all addresses, use the IP address **0.0.0.0** with a network mask of **0.0.0.0**.
- Step 14** From the Netmask drop-down list, choose the network mask applied to the network IP address. You can type a network mask into this field or select one of the common masks from the list.

Step 15 Click **OK**.

Customize 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:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > EIGRP > Setup . The EIGRP Setup pane appears.
Step 2	Check the Enable EIGRP routing check box.
Step 3	Click OK .
Step 4	Choose Configuration > Device Setup > Routing > EIGRP > Interfaces . The Interface pane appears and displays all of the EIGRP interface configurations.
Step 5	Double-click an interface entry or choose the entry and click Edit . The Edit EIGRP Interface Entry dialog box appears.
Step 6	Choose the EIGRP AS number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.
Step 7	In the Hello Interval field, enter the interval between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Step 8	In the Hold Time field, specify the hold time, in seconds. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
Step 9	Click OK .
-

Disable 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.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**. The EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** Click the **Process Instance** tab.
- Step 4** Click **Advanced**.
- Step 5** In the Summary area, uncheck the **Auto-Summary** check box.



Note This setting is enabled by default.

- Step 6** Click **OK**.
-

Configure 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.

Procedure

In ASDM, the Default Information pane displays a table of rules for controlling the sending and receiving of default route information in EIGRP updates. You can have one in and one out rule for each EIGRP routing process (only one process is currently supported).

By default, default routes are sent and accepted. To restrict or disable the sending and receiving of default route information, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**. The main EIGRP Setup pane appears.
- Step 2** Check the **Enable EIGRP routing** check box.
- Step 3** Click **OK**.
- Step 4** Do one of the following:
- Click **Add** to create a new entry.
 - To edit an entry, double-click the entry in the table or select an entry in the table and click **Edit**.

The Add Default Information or Edit Default Information dialog box appears for that entry. The EIGRP AS number is automatically selected in the EIGRP field.

- Step 5** In the Direction field, choose the direction for the rule from the following options:
- **in**—The rule filters default route information from incoming EIGRP updates.
 - **out**—The rule filters default route information from outgoing EIGRP updates.
- You can have one in rule and one out rule for each EIGRP process.
- Step 6** Add network rules to the network rule table. The network rules define which networks are allowed and which are not when receiving or sending default route information. Repeat the following steps for each network rule you are adding to the default information filter rule.
- a. Click **Add** to add a network rule. Double-click an existing network rule to edit the rule.
 - b. In the Action field, click **Permit** to allow the network or **Deny** to block the network.
 - c. Enter the IP address and network mask of the network being permitted or denied by the rule in the IP Address and Network Mask fields.

To deny all default route information from being accepted or sent, enter **0.0.0.0** as the network address and choose **0.0.0.0** as the network mask.
 - d. Click **OK** to add the specified network rule to the default information filter rule.
- Step 7** Click **OK** to accept the default information filter rule.
-

Disable 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Interfaces**. The Interface pane appears and displays the EIGRP interface configurations.
- Step 2** Double-click an interface entry or choose the entry and click **Edit**. The Edit EIGRP Interface Entry dialog box appears.
- Step 3** Choose the EIGRP Autonomous system (AS) number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.
- Step 4** Uncheck the **Split Horizon** check box.

Step 5 Click **OK**.

Restart the EIGRP Process

You can restart an EIGRP process or clear redistribution or clear counters.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > EIGRP > Setup**.
The EIGRP Setup pane appears.
- Step 2** Click **Reset**.
-

Monitoring for EIGRP

You can use the following commands to monitor the EIGRP routing process. For examples and descriptions of the command output, see the command reference. Additionally, you can disable the logging of neighbor change messages and neighbor warning messages.

To monitor or disable various EIGRP routing statistics, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Monitoring > Routing > EIGRP Neighbor**.
- Each row represents one EIGRP neighbor. For each neighbor, the list includes its IP address, the interface to which the neighbor is connected, the holdtime, the uptime, the queue length, the sequence number, the smoothed round trip time, and the retransmission timeout. The list of possible state changes are the following:
- **NEW ADJACENCY**—A new neighbor has been established.
 - **PEER RESTARTED**—The other neighbor initiates the reset of the neighbor relationship. The router getting the message is not the one resetting the neighbor.
 - **HOLD TIME EXPIRED**—The router has not heard any EIGRP packets from the neighbor within the hold-time limit.
 - **RETRY LIMIT EXCEEDED**—EIGRP did not receive the acknowledgement from the neighbor for EIGRP reliable packets, and EIGRP has already tried to retransmit the reliable packet 16 times without any success.
 - **ROUTE FILTER CHANGED**—The EIGRP neighbor is resetting because there is a change in the route filter.
 - **INTERFACE DELAY CHANGED**—The EIGRP neighbor is resetting because there is a manual configuration change in the delay parameter on the interface.
 - **INTERFACE BANDWIDTH CHANGED**—The EIGRP neighbor is resetting because there is a manual configuration change in the interface bandwidth on the interface.
 - **STUCK IN ACTIVE**—The EIGRP neighbor is resetting because EIGRP is stuck in active state. The neighbor getting reset is the result of the stuck-in-active state.

- Step 2** Click the EIGRP neighbor that you want to monitor.
- Step 3** To remove the current list of neighbors, click **Clear Neighbors**.
- Step 4** To refresh the current list of neighbors, click **Refresh**.

**Note**

By default, neighbor change and neighbor warning messages are logged.

History for EIGRP

Table 27-1 *Feature History for EIGRP*

Feature Name	Platform Releases	Feature Information
EIGRP Support	7.0(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP). We introduced the following screen: Configuration > Device Setup > Routing > EIGRP.
Dynamic Routing in Multiple Context Mode	9.0(1)	EIGRP routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > EIGRP > Setup.
Clustering	9.0(1)	For EIGRP, bulk synchronization, route synchronization, and layer 2 load balancing are supported in the clustering environment.
EIGRP Auto-Summary	9.2(1)	For EIGRP, the Auto-Summary field is now disabled by default. We modified the following screen: Configuration > Device Setup > Routing > EIGRP > Setup > Edit EIGRP Process Advanced Properties



Multicast Routing

This chapter describes how to configure the Cisco ASA to use the multicast routing protocol.

- [About Multicast Routing, page 28-1](#)
- [Guidelines for Multicast Routing, page 28-3](#)
- [Enable Multicast Routing, page 28-3](#)
- [Customize Multicast Routing, page 28-4](#)
- [Example for Multicast Routing, page 28-18](#)
- [History for Multicast Routing, page 28-19](#)

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.

- [Stub Multicast Routing, page 28-2](#)
- [PIM Multicast Routing, page 28-2](#)
- [Multicast Group Concept, page 28-2](#)
- [Clustering, page 28-2](#)

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. For information about how to configure multicast groups, see [Configure a Multicast Group](#), page 28-14.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Layer 2 clustering, the master unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, slave units may forward multicast data packets. All data flows are full flows. Stub forwarding flows are also supported. Because only one unit receives multicast packets in Layer 2

clustering, redirection to the master unit is common. In Layer 3 clustering, units do not act independently. All data and routing packets are processed and forwarded by the master unit. Slave units drop all packets that have been sent.

For more information about clustering, see [Chapter 8, “ASA Cluster.”](#)

Guidelines for Multicast Routing

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.

Enable Multicast Routing

Enabling multicast routing on the ASA, 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.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast** routing check box.

Checking this check box enables IP multicast routing on the ASA. Unchecking this check box disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces. You can disable multicast on a per-interface basis.

Table 28-1 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 28-1 **Entry Limits for Multicast Tables**

Table	16 MB	128 MB	128+ MB
MFIB	1000	3000	30000
IGMP Groups	1000	3000	30000
PIM Routes	3000	7000	72000

Customize Multicast Routing

This section describes how to customize multicast routing.

- [Configure Stub Multicast Routing and Forward IGMP Messages, page 28-4](#)
- [Configure a Static Multicast Route, page 28-5](#)
- [Configure IGMP Features, page 28-6](#)
- [Configure PIM Features, page 28-10](#)
- [Configure a Multicast Group, page 28-14](#)
- [Configure a Bidirectional Neighbor Filter, page 28-16](#)
- [Configure a Multicast Boundary, page 28-17](#)

Configure Stub Multicast Routing and Forward 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.

Procedure

- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.
- Step 2** In the Multicast pane, check the **Enable Multicast routing** check box.
- Step 3** Click **Apply** to save your changes.
- Step 4** Choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.
- Step 5** To modify the specific interface from which you want to forward IGMP messages, select the interface and click **Edit**.

The Configure IGMP Parameters dialog box appears.

- Step 6** From the **Forward Interface** drop-down list, choose the specific interface from which you want to forward IGMP messages.
- Step 7** Click **OK** to close this dialog box, then click **Apply** to save your changes.
-

Configure 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.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > MRoute**.
- Step 2** Choose **Add** or **Edit**.
- The Add or Edit Multicast Route dialog box appears.
- Use the Add Multicast Route dialog box to add a new static multicast route to the ASA. Use the Edit Multicast Route dialog box to change an existing static multicast route.
- Step 3** In the Source Address field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.
- Step 4** Choose the network mask for the IP address of the multicast source from the Source Mask drop-down list.
- Step 5** In the Incoming Interface area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:
- In the Source Interface field, choose the incoming interface for the multicast route from the drop-down list.
 - In the Destination Interface field, choose the destination interface that the route is forwarded through from the drop-down list.
-  **Note** You can specify the interface or the RPF neighbor, but not both at the same time.
-
- Step 6** In the Administrative Distance field, choose the administrative distance of the static multicast route. If the static multicast route has the same administrative distance as the unicast route, then the static multicast route takes precedence.
- Step 7** Click **OK**.
-

Configure 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.

- [Disable IGMP on an Interface, page 28-6](#)
- [Configure IGMP Group Membership, page 28-7](#)
- [Configure a Statically Joined IGMP Group, page 28-7](#)
- [Control Access to Multicast Groups, page 28-8](#)
- [Limit the Number of IGMP States on an Interface, page 28-9](#)
- [Modify the Query Messages to Multicast Groups, page 28-9](#)
- [Change the IGMP Version, page 28-10](#)

Disable 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.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.
- The Protocol pane displays the IGMP parameters for each interface on the ASA.
- Step 2** Choose the interface that you want to disable and click **Edit**.
- Step 3** To disable the specified interface, uncheck the **Enable IGMP** check box.
- Step 4** Click **OK**.

The Protocol pane displays Yes if IGMP is enabled on the interface, or No if IGMP is disabled on the interface.

Configure 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 [Configure a Statically Joined IGMP Group, page 28-7](#).

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Join Group**.
 - Step 2** Click **Add** or **Edit** in the **Join Group** pane.

The **Add IGMP Join Group** dialog box allows you to configure an interface to be a member of a multicast group. The **Edit IGMP Join Group** dialog allows you to change existing membership information.
 - Step 3** In the **Interface Name** field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.
 - Step 4** In the **Multicast Group Address** field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.
 - Step 5** Click **OK**.
-

Configure a Statically Joined IGMP Group

Sometimes a group member cannot report its membership in the group because of some configuration, or there may be no members of a group on the network segment. However, you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group.

In the main ASDM window, choose **Configuration > Routing > Multicast > IGMP > Static Group** to configure the ASA to be a statically connected member of a group. With this method, the ASA does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but this interface is not a member of the multicast group.

Procedure:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Static Group**.
 - Step 2** Click **Add** or **Edit** in the **Static Group** pane.

Use the **Add IGMP Static Group** dialog box to statically assign a multicast group to an interface. Use the **Edit IGMP Static Group** dialog box to change existing static group assignments.

- Step 3** In the **Interface Name** field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.
- Step 4** In the **Multicast Group Address** field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.
- Step 5** Click **OK**.
-

Control Access to Multicast Groups

You can control access to multicast groups by using access control lists.

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Access Group**.
- The **Access Group** pane appears. The table entries in the Access Group pane are processed from the top down. Place more specific entries near the top of the table and more generic entries further down. For example, place an access group entry that permits a specific multicast group near the top of the table and an access group entry below that denies a range of multicast groups, including the group in the permit rule. The group is permitted because the permit rule is enforced before the deny rule.
- Double-clicking an entry in the table opens the [Add or Edit Access Group](#) dialog box for the selected entry.
- Step 2** Click **Add** or **Edit**.
- The **Add Access Group** or **Edit Access Group** dialog box appears. The **Add Access Group** dialog box lets you add a new access group to the Access Group Table. The **Edit Access Group** dialog box lets you change information for an existing access group entry. Some fields may be dimmed when editing existing entries.
- Step 3** Choose the interface name with which the access group is associated from the **Interface** drop-down list. You cannot change the associated interface when you are editing an existing access group.
- Step 4** Choose **permit** from the **Action** drop-down list to allow the multicast group on the selected interface. Choose **deny** from the **Action** drop-down list to filter the multicast group from the selected interface.
- Step 5** In the **Multicast Group Address** field, enter the address of the multicast group to which the access group applies.
- Step 6** Enter the network mask for the multicast group address, or choose one of the common network masks from the **Netmask** drop-down list.
- Step 7** Click **OK**.
-

Limit 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.

Procedure:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.
- Step 2** Choose the interface you want to limit from the table on the **Protocol** pane, and click **Edit**.
The **Configure IGMP Parameters** dialog box appears.
- Step 3** Enter the maximum number of host that can join on an interface, in the **Group Limit** field.
The default value is 500. Valid values range from 0 to 500.
-  **Note** Setting this value to 0 prevents learned groups from being added, but manually defined memberships are still permitted.
-
- Step 4** Click **OK**.
-

Modify the Query Messages to Multicast Groups

The ASA sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the ASA. If the ASA discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packet for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the ASA does not receive a response to a host query within this amount of time, it deletes the group.

To change the query interval, query response time, and query timeout value, perform the following steps:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.
- Step 2** Choose the interface you want to limit from the table on the **Protocol** pane, and click **Edit**.
The **Configure IGMP Parameters** dialog box appears.

- Step 3** Enter the interval in seconds, at which the designated router sends IGMP host-query messages, in the **Query Interval** field.

Valid values range from 1 to 3600 seconds. The default value is 125 seconds.



Note If the ASA does not hear a query message on an interface for the specified timeout value, then the ASA becomes the designated router and starts sending the query messages.

- Step 4** Enter the period of time, in seconds, in the **Query Timeout** field before which the ASA takes over as the requester for the interface after the previous requester has stopped doing so.

Valid values range from 60 to 300 seconds. The default value is 255 seconds.

- Step 5** Click **OK**.

Change the IGMP Version

By default, the ASA runs IGMP Version 2, which enables several additional features.

All multicast routers on a subnet must support the same version of IGMP. The ASA does not automatically detect Version 1 routers and switch to Version 1. However, a mix of IGMP Version 1 and 2 hosts on the subnet works; the ASA running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

To control which version of IGMP is running on an interface, perform the following steps.

Procedure

- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Protocol**.

- Step 2** Choose the interface whose version of IGMP you want to change from the table on the **Protocol** pane, and click **Edit**.

The **Configure IGMP Interface** dialog box appears.

- Step 3** Choose the version number from the **Version** drop-down list.

- Step 4** Click **OK**.

Configure 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.

- [Enable and Disable PIM on an Interface, page 28-11](#)
- [Configure a Static Rendezvous Point Address, page 28-11](#)
- [Configure the Designated Router Priority, page 28-12](#)
- [Configure and Filter PIM Register Messages, page 28-13](#)
- [Configure PIM Message Intervals, page 28-13](#)
- [Configure a Route Tree, page 28-14](#)
- [Filter PIM Neighbors, page 28-15](#)

Enable and Disable 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:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > Multicast > PIM > Protocol .
Step 2	Choose the interface on which you want to enable PIM from the table on the Protocol pane, and click Edit . The Edit PIM Protocol dialog box appears.
Step 3	Check the Enable PIM check box. To disable PIM, uncheck this check box.
Step 4	Click OK .
-

Configure a Static Rendezvous Point Address

All routers within a common PIM sparse mode or bidir domain require knowledge of the PIM RP address. The address is statically configured using the **pim rp-address** command.



Note

The ASA does not support Auto-RP or PIM BSR

You can configure the ASA to serve as RP to more than one group. The group range specified in the ACL determines the PIM RP group mapping. If an ACL is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4).

To configure the address of the PIM PR perform the following steps.

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > Multicast > PIM > Rendezvous Points .
Step 2	Click Add or Edit .

The **Add or Edit Rendezvous Point** dialog box appears. The **Add Rendezvous Point** dialog box lets you add a new entry to the Rendezvous Point table. The **Edit Rendezvous Point** dialog box lets you change an existing RP entry. Additionally, you can click **Delete** to remove the selected multicast group entry from the table.

These restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Step 3 In the **Rendezvous Point Address** field, enter the IP address for the RP.

When editing an existing RP entry, you cannot change this value.

Step 4 Check the **Use bi-directional forwarding** check box if the specified multicast groups are to operate in bidirectional mode. The Rendezvous Point pane displays Yes if the specified multicast groups are to operate in bidirectional mode and displays No if the specified groups are to operate in sparse mode. In bidirectional mode, if the ASA receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.

Step 5 Click the **Use this RP for All Multicast Groups** radio button to use the specified RP for all multicast groups on the interface, or the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.

For more information about multicast groups, see [Configure a Multicast Group, page 28-14](#).

Step 6 Click **OK**.

Configure the Designated Router Priority

The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, selecting the DR is based on the DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR.

By default, the ASA has a DR priority of 1. To change this value perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Protocol**.

Step 2 Choose the interface that you want to enable for PIM from the table on the **Protocol** pane, and click **Edit**. The **Edit PIM Protocol** dialog box appears.

Step 3 In the **DR Priority** field, type the value for the designated router priority for the selected interface. The router with the highest DR priority on the subnet becomes the designated router. Valid values range from 0 to 4294967294. The default DR priority is 1. Setting this value to 0 makes the ASA interface ineligible to become the default router.

Step 4 Click **OK**.

Configure and Filter PIM Register Messages

When the ASA is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. The Request Filter pane lets you define the multicast sources from which the ASA will accept PIM register messages.

To filter PIM register messages perform the following steps:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > Multicast > PIM > Request Filter .
Step 2	Click Add . The Request Filter Entry dialog box lets you define the multicast sources that are allowed to register with the ASA when the ASA acts as an RP. You create the filter rules based on the source IP address and the destination multicast address.
Step 3	From the Action drop-down list, choose Permit to create a rule that allows the specified source of the specified multicast traffic to register with the ASA, or choose Deny to create a rule that prevents the specified source of the specified multicast traffic from registering with the ASA.
Step 4	Type the IP address for the source of the register message, in the Source IP Address field.
Step 5	Type or choose the network mask from the drop-down list for the source of the register message, in the Source Netmask field.
Step 6	Type the multicast destination address, in the Destination IP Address field.
Step 7	Type or choose the network mask from the drop-down list for the multicast destination address, in the Destination Netmask field.
Step 8	Click OK .
-

Configure 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:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > Multicast > PIM > Protocol .
Step 2	Choose the interface that you want to enable for PIM from the table on the Protocol pane, and click Edit . The Edit PIM Protocol dialog box appears.
Step 3	Type the frequency, in seconds, at which the interface sends PIM hello messages, in the Hello Interval field.

- Step 4** Type the frequency, in seconds, at which the interface sends PIM join and prune advertisements, in the **Prune Interval** field.
- Step 5** Click **OK**.
-

Configure a Route Tree

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the ASA should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

To configure a PIM leaf router tree, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Route Tree**.
- Step 2** Click one of the following radio buttons:
- **Use Shortest Path Tree for All Groups**—Choose this option to use the shortest-path tree for all multicast groups.
 - **Use Shared Tree for All Groups**—Choose this option to use the shared tree for all multicast groups.
 - **Use Shared Tree for the Groups specified below**—Choose this option to use the shared tree for the groups specified in the Multicast Groups table. The shortest-path tree is used for any group that is not specified in the Multicast Groups table.

The **Multicast Groups** table displays the multicast groups to use with the shared tree.

The table entries are processed from the top down. You can create an entry that includes a range of multicast groups, but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

To edit a multicast group, see [Configure a Multicast Group, page 28-14](#).

Configure a Multicast Group

Multicast groups are lists of access rules that define which multicast addresses are part of a group. A multicast group can include a single multicast address or a range of multicast addresses. Use the Add Multicast Group dialog box to create a new multicast group rule. Use the Edit Multicast Group dialog box to modify an existing multicast group rule.

To configure a multicast group, perform the following steps:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Rendezvous Points**.
- Step 2** The **Rendezvous Point** pane appears. Click the group that you want to configure.
- The **Edit Rendezvous Point** dialog box appears.

- Step 3** Click the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.
- Step 4** Click **Add** or **Edit**.
The **Add or Edit Multicast Group** dialog box appears.
- Step 5** From the **Action** drop-down list, choose **Permit** to create a group rule that allows the specified multicast addresses, or choose **Deny** to create a group rule that filters the specified multicast addresses.
- Step 6** In the Multicast Group Address field, type the multicast address associated with the group.
- Step 7** From the Netmask drop-down list, choose the network mask for the multicast group address.
- Step 8** Click **OK**.
-

Filter 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:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > PIM > Neighbor Filter**.
- Step 2** Choose the PIM neighbor that you want to configure from the table by clicking **Add/Edit/Insert**.
The **Add/Edit/Insert Neighbor Filter** Entry dialog box appears. It lets you create the ACL entries for the multicast boundary ACL. You can also delete a selected PIM neighbor entry.
- Step 3** Choose the interface name from the **Interface Name** drop-down list.
- Step 4** From the **Action** drop-down list, choose **Permit** or **Deny** for the neighbor filter ACL entry.
Choosing **Permit** allows the multicast group advertisements to pass through the interface. Choosing **Deny** prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.
- Step 5** Enter the IP address of the multicast PIM group being permitted or denied, in the **IP Address** field. Valid group addresses range from 224.0.0.0 to 239.255.255.255.
- Step 6** From the **Netmask** drop-down list, choose the netmask for the multicast group address.
- Step 7** Click **OK**.
-

Configure 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:

Procedure

-
Step 1	In the main ASDM window, choose Configuration > Device Setup > Routing > Multicast > PIM > Bidirectional Neighbor Filter .
Step 2	Double-click an entry in the PIM Bidirectional Neighbor Filter table to access the Edit Bidirectional Neighbor Filter Entry dialog box for that entry.
Step 3	Choose the PIM neighbor that you want to configure from the table by clicking Add/Edit/Insert . The Add/Edit/Insert Bidirectional Neighbor Filter Entry dialog box appears, which lets you create ACL entries for the PIM bidirectional neighbor filter ACL.
Step 4	Choose the interface name from the Interface Name drop-down list. Select the interface for which you are configuring the PIM bidirectional neighbor filter ACL entry.
Step 5	From the Action drop-down list, choose Permit or Deny for the neighbor filter ACL entry. Choose Permit to allow the specified devices to participate in the DF election process. Choose Deny to prevent the specified devices from participating in the DF election process.
Step 6	Enter the IP address of the multicast PIM group being permitted or denied. Valid group addresses range from 224.0.0.0 to 239.255.255.255, in the IP Address field.
Step 7	From the Netmask drop-down list, choose the netmask for the multicast group address.

Step 8 Click **OK**.

Configure a Multicast Boundary

Address scoping defines domain boundaries so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary on an interface for multicast group addresses by choosing **Configuration > Routing > Multicast > MBoundary** in ASDM. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

A standard ACL defines the range of affected addresses. When a boundary is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

To configure a multicast boundary perform the following steps:

Procedure

-
- Step 1** In the main ASDM window, choose **Configuration > Routing > Multicast > MBoundary**.
- The MBoundary pane lets you configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.
- Step 2** Click **Edit**.
- The **Edit Boundary Filter** dialog box appears and displays the multicast boundary filter ACL. You can add and remove boundary filter ACL entries using this dialog box.
- When the boundary filter configuration is applied to the ASA, the ACL appears in the running configuration with the name *interface-name_multicast*, where the *interface-name* is the name of the interface to which the multicast boundary filter is applied. If an ACL with that name already exists, a number is appended to the name (for example, *inside_multicast_1*).
- Step 3** Choose the interface for which you are configuring the multicast boundary filter ACL from the **Interface** drop-down list.
- Step 4** Check the **Remove any Auto-RP group range** check box to filter Auto-RP messages from sources denied by the boundary ACL. If the **Remove any Auto-RP group range** check box is unchecked, all Auto-RP messages are passed.

Step 5 Click **OK**.

Example for Multicast Routing

The following example shows how to enable and configure multicast routing with various optional processes:

-
- Step 1** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast**.
- Step 2** In the Multicast pane, check the **Enable Multicast** routing check box, and click **Apply**.
- Step 3** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > MRoute**.
- Step 4** Click **Add** or **Edit**.

The **Add or Edit Multicast Route** dialog box appears.

Use the **Add Multicast Route** dialog box to add a new static multicast route to the ASA. Use the **Edit Multicast Route** dialog box to change an existing static multicast route.

- Step 5** In the **Source Address** field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.
- Step 6** Choose the network mask for the IP address of the multicast source from the **Source Mask** drop-down list.
- Step 7** In the **Incoming Interface** area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:
- In the **Source Interface** field, choose the incoming interface for the multicast route from the drop-down list.
 - In the **Destination Interface** field, choose the destination interface to which the route is forwarded through the selected interface from the drop-down list.



Note You can specify the interface or the RPF neighbor, but not both at the same time.

- Step 8** In the **Administrative Distance** field, choose the administrative distance of the static multicast route. If the static multicast route has the same administrative distance as the unicast route, then the static multicast route takes precedence.
- Step 9** Click **OK**.
- Step 10** In the main ASDM window, choose **Configuration > Device Setup > Routing > Multicast > IGMP > Join Group**.
- The **Join Group** pane appears.
- Step 11** Click **Add** or **Edit**.
- The **Add IGMP Join Group** dialog box allows you to configure an interface to be a member of a multicast group. The **Edit IGMP Join Group** dialog box allows you to change existing membership information.
- Step 12** In the **Interface Name** field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

- Step 13** In the **Multicast Group Address** field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.
- Step 14** Click **OK**.
-

History for Multicast Routing

Table 28-2 Feature History for Multicast Routing

Feature Name	Platform Releases	Feature Information
Multicast routing support	7.0(1)	Support was added for multicast routing data, authentication, and redistribution and monitoring of routing information using the multicast routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > Multicast.
Clustering support	9.0(1)	Support was added for clustering.



IPv6 Neighbor Discovery

- [About IPv6 Neighbor Discovery, page 29-1](#)
- [Prerequisites for IPv6 Neighbor Discovery, page 29-5](#)
- [Guidelines for IPv6 Neighbor Discovery, page 29-5](#)
- [Defaults for IPv6 Neighbor Discovery, page 29-6](#)
- [Configure IPv6 Neighbor Discovery, page 29-7](#)
- [View and Clear Dynamically Discovered Neighbors, page 29-12](#)
- [History for IPv6 Neighbor Discovery, page 29-13](#)

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.

- [Neighbor Solicitation Messages, page 29-2](#)
- [Neighbor Reachable Time, page 29-3](#)
- [Duplicate Address Detection, page 29-3](#)
- [Router Advertisement Messages, page 29-3](#)
- [Static IPv6 Neighbors, page 29-5](#)


Neighbor Solicitation Messages

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. The neighbor solicitation message is sent to the solicited-node multicast address. The source address in the neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The neighbor solicitation message also includes the link-layer address of the source node.

After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (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. Figure 29-1 shows the neighbor solicitation and response process.

Figure 29-1 IPv6 Neighbor Discovery—Neighbor Solicitation Message



Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to 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 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

The Cisco ASA can participate in router advertisements so that neighboring devices can dynamically learn a default router address. Router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface of the ASA. The router advertisement messages are sent to the all-nodes multicast address. [Figure 29-2](#) shows an example of how to send router advertisement messages on an IPv6 configured interface.

Figure 29-2 IPv6 Neighbor Discovery—Router Advertisement Message

Router advertisement messages typically include the following information:

- One or more IPv6 prefix that nodes on the local link can use to automatically configure their IPv6 addresses.
- Lifetime information for each prefix included in the advertisement.
- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed.
- Default router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router).
- Additional information for hosts, such as the hop limit and MTU a host should use in packets that it originates.
- The amount of time between neighbor solicitation message retransmissions on a given link.
- The amount of time a node considers a neighbor reachable.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message. Because router solicitation messages are usually sent by hosts at system startup, and the host does not have a configured unicast address, the source address in router solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the router solicitation message is used as the source address in the message. The destination address in router solicitation messages is the all-routers multicast address with a scope of the link. When a router advertisement is sent in response to a router solicitation, the destination address in the router advertisement message is the unicast address of the source of the router solicitation message.

You can configure the following settings for router advertisement messages:

- The time interval between periodic router advertisement messages.
- The router lifetime value, which indicates the amount of time IPv6 nodes should consider the ASA to be the default router.
- The IPv6 network prefixes in use on the link.
- Whether or not an interface transmits router advertisement messages.

Unless otherwise noted, the router advertisement message settings are specific to an interface and are entered in interface configuration mode.

Static IPv6 Neighbors

You can manually define a neighbor in the IPv6 neighbor cache. If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry. Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.

Prerequisites for IPv6 Neighbor Discovery

Configure IPv6 addressing according to [Configure IPv6 Addressing, page 15-12](#).

Guidelines for IPv6 Neighbor Discovery

Firewall Mode Guidelines

The following IPv6 neighbor discovery commands are not supported in transparent firewall mode, because they require router capabilities:

- **ipv6 nd prefix**
- **ipv6 nd ra-interval**
- **ipv6 nd ra-lifetime**
- **ipv6 nd suppress-ra**

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.

Defaults for IPv6 Neighbor Discovery

Table 29-1 lists the default settings for IPv6 Neighbor Discovery.

Table 29-1 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.

Table 29-1 Default IPv6 Neighbor Discovery Parameters (continued)

Parameters	Default
<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.

Configure IPv6 Neighbor Discovery

- [Configure the Neighbor Solicitation Message Interval, page 29-7](#)
- [Configure the Neighbor Reachable Time, page 29-8](#)
- [Configure the Router Advertisement Transmission Interval, page 29-8](#)
- [Configure the Router Lifetime Value, page 29-9](#)
- [Configure DAD Settings, page 29-9](#)
- [Suppress Router Advertisement Messages, page 29-10](#)
- [Configure Address Config Flags for IPv6 DHCP Relay, page 29-10](#)
- [Configure the IPv6 Prefix in Router Advertisements, page 29-11](#)
- [Configuring a Static IPv6 Neighbor, page 29-12](#)

Configure the Neighbor Solicitation Message Interval

To configure the interval between IPv6 neighbor solicitation retransmissions on an interface perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Choose the interface on which to configure the neighbor solicitation interval. The interface must have been configured with an IPv6 address. See [Configure IPv6 Addressing, page 15-12](#) for more information.
- Step 3** Click **Edit**. The **Edit Interface** dialog box appears.

- Step 4** Click the **IPv6** tab.
 - Step 5** Enter the time interval, in the **NS Interval** field.
 - Step 6** Click **OK**.
 - Step 7** Click **Apply** to save the running configuration.
-

Configure the Neighbor Reachable Time

To configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
 - Step 2** Choose the interface for which you want to configure the time. The interface must have been configured with an IPv6 address. For more information, see [Configure IPv6 Addressing, page 15-12](#).
 - Step 3** Click **Edit**. The **Edit Interface** dialog box appears.
 - Step 4** Click the **IPv6** tab.
 - Step 5** Enter a valid value in the **Reachable Time** field.
 - Step 6** Click **OK**.
 - Step 7** Click **Apply** to save the running configuration.
-

Configure the Router Advertisement Transmission Interval

To configure the interval between IPv6 router advertisement transmissions on an interface, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
 - Step 2** Select the interface for which you want to configure the time.
The interface must have been configured with an IPv6 address. For more information, see [Configure IPv6 Addressing, page 15-12](#).
 - Step 3** Click **Edit**. The **Edit Interface** dialog box appears.
 - Step 4** Click the **IPv6** tab.
 - Step 5** Enter a valid transmission interval value, in the **RA Interval** field,.

**Note**

(Optional) To add a router advertisement transmission interval value in milliseconds instead, check the **RA Interval in Milliseconds** check box, and enter a value from 500 to 1800000.

- Step 6** Click **OK**.
- Step 7** Click **Apply** to save the running configuration.
-

Configure the Router Lifetime Value

To configure the router lifetime value in IPv6 router advertisements on an interface perform the following steps.

Procedure

- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select the interface you want to configure.
The interface must have been configured with an IPv6 address. For more information see [Configure IPv6 Addressing, page 15-12](#).
- Step 3** Click **Edit**.
The **Edit Interface** dialog box appears.
- Step 4** Click the **IPv6** tab.
- Step 5** Enter a valid lifetime value, in the **RA Lifetime** field.
- Step 6** Click **OK**.
- Step 7** Click **Apply** to save the running configuration.
-

Configure DAD Settings

To specify DAD settings on the interface, perform the following steps.

Procedure

- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select the interface you want to configure.
The interface must have been configured with an IPv6 address. For more information, see [Configure IPv6 Addressing, page 15-12](#).
- Step 3** Click **Edit**.
The **Edit Interface** dialog box appears.
- Step 4** Click the **IPv6** tab.

Step 5 Enter the number of allowed DAD attempts.

This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses.

- Valid values range from 0 to 600.
 - A zero value disables DAD processing on the specified interface. The default is one message.
-

Suppress Router Advertisement Messages

Router advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want the ASA to supply the IPv6 prefix (for example, the outside interface).

To suppress the router lifetime value in IPv6 router advertisements on an interface, perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select the interface for which you want to suppress the router advertisement transmissions. The interface must have been configured with an IPv6 address.
- Step 3** Click **Edit**.
The **Edit Interface** dialog box appears.
- Step 4** Click the **IPv6** tab.
- Step 5** Check the **Suppress RA** check box.
-

Configure 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.

Procedure

-
- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select the interface that you want to configure.
- Step 3** Click **Edit**.
The **Edit Interface** dialog box appears.
- Step 4** Click the **IPv6** tab.
- Step 5** Check the **Hosts should use DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet.

This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

- Step 6** Check the **Hosts should use DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Configure the IPv6 Prefix in Router Advertisements

To configure the which IPv6 prefixes are included in IPv6 router advertisements, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Setup > Interface Settings > Interfaces**.
- Step 2** Select the interface for which you want to suppress the router advertisement transmissions. The interface must have been configured with an IPv6 address.
- Step 3** Click **Edit**.
The **Edit Interface** dialog box appears.
- Step 4** Click the **IPv6** tab.
- Step 5** In the **Interface IPv6 Prefixes** area, click **Add**.
The **Add IPv6 Prefix for Interface** dialog box appears.
- Step 6** Enter the IPv6 address with the prefix length.
- Step 7** (Optional) To configure the IPv6 address manually, check the **No Auto-Configuration** check box. This setting indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.
- Step 8** (Optional) To indicate that the IPv6 prefix is not advertised, check the **No Advertisements** check box.
- Step 9** (Optional) The **Off Link** check box indicates that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.
- Step 10** In the **Prefix Lifetime** area, click the **Lifetime Duration** radio button, and specify the following:
- A valid lifetime for the prefix in seconds from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days).
 - A preferred lifetime for the prefix from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days).
- Step 11** To define a prefix lifetime expiration date, click the **Lifetime Expiration Date** radio button, and specify the following:
- Choose a valid month and day from the drop-down list, and then enter a time in hh:mm format.
 - Choose a preferred month and day from the drop-down list, and then enter a time in hh:mm format.

Step 12 Click **OK** to save your settings.

The **Interface IPv6 Prefixes Address** field appears with the preferred and valid dates.

Configuring a Static IPv6 Neighbor

Make sure that IPv6 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed.

For information about configuring IPv6 addresses, see [Configure IPv6 Addressing, page 15-12](#).

To add an IPv6 static neighbor, perform the following steps.

Procedure

Step 1 Choose **Configuration > Device Management > Advanced > IPv6 Neighbor Discovery Cache**.

Step 2 Click **Add**.

The Add IPv6 Static Neighbor dialog box appears.

Step 3 From the Interface Name drop-down list, choose an interface on which to add the neighbor.

Step 4 In the IP Address field, enter the IPv6 address that corresponds to the local data-link address, or click the ellipsis (...) to browse for an address.

If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry.

Step 5 In the MAC address field, enter the local data-line (hardware) MAC address.

Step 6 Click **OK**.



Note Before you apply the changes and save the configuration, you can click **Reset** to cancel any changes and restore the original values.

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

View and Clear Dynamically Discovered Neighbors

When a host or node communicates with a neighbor, the neighbor is added to the neighbor discovery cache. The neighbor is removed from the cache when there is no longer any communication with that neighbor.

To view dynamically discovered neighbors and clear these neighbors from the IPv6 neighbor discovery cache, perform the following steps:

Step 1 Choose **Monitoring > Interfaces > IPv6 Neighbor Discovery Cache**.

You can view all static and dynamically discovered neighbors from the IPv6 Neighbor Discovery Cache pane.

- Step 2** To clear all dynamically discovered neighbors from the cache, click **Clear Dynamic Neighbor Entries**. The dynamically discovered neighbor is removed from the cache.



Note This procedure clears only dynamically discovered neighbors from the cache; it does not clear static neighbors.

History for IPv6 Neighbor Discovery

Table 29-2 Feature History for IPv6 Neighbor Discovery

Feature Name	Releases	Feature Information
IPv6 Neighbor Discovery	7.0(1)	We introduced this feature. We introduced the following screens: Monitoring > Interfaces > IPv6 Neighbor Discovery Cache. Configuration > Device Management > Advanced > IPv6 Neighbor Discovery Cache. Configuration > Device Setup > Interface Settings > Interfaces > IPv6.
Address Config Flags for IPv6 DHCP Relay	9.0(1)	We modified the following screen: Configuration > Device Device Setup > Interfaces > IPv6.



PART 6

AAA Servers and the Local Database



AAA and the Local Database

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 also describes how to configure the local database for AAA functionality. For external AAA servers, see the chapter for your server type.

- [About AAA and the Local Database, page 30-1](#)
- [Guidelines for the Local Database, page 30-4](#)
- [Add a User Account to the Local Database, page 30-5](#)
- [Generate a Shared Key, page 30-6](#)
- [Test Local Database Authentication and Authorization, page 30-8](#)
- [Monitoring the Local Database, page 30-9](#)
- [History for the Local Database, page 30-9](#)

About AAA and the Local Database

This section describes AAA and the local database.

- [Authentication, page 30-2](#)
- [Authorization, page 30-2](#)
- [Accounting, page 30-2](#)
- [Interaction Between Authentication, Authorization, and Accounting, page 30-2](#)
- [AAA Servers, page 30-3](#)
- [AAA Server Groups, page 30-3](#)
- [About the Local Database, page 30-3](#)

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 Cisco ASA to authenticate the following items:

- All administrative connections to the ASA, including the following sessions:
 - Telnet
 - SSH
 - Serial console
 - ASDM using HTTPS
 - VPN management access
- The **enable** command
- Network access
- VPN access

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
- Network access
- VPN access

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.

About the Local Database

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.

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 Cisco ASDM login.

- Command authorization

If you turn on command authorization using the local database, then the Cisco ASA refers to the user privilege level to determine which commands are available. Otherwise, the privilege level is not generally used. By default, all commands are either privilege level 0 or level 15. ASDM allows you to enable three predefined privilege levels, with commands assigned to level 15 (Admin), level 5 (Read Only), and level 3 (Monitor Only). If you use the predefined levels, then assign users to one of these three privilege levels.

- Network access authentication
- VPN client authentication

For multiple context mode, you can configure usernames in the system execution space to provide individual logins at the CLI using the **login** command; however, you cannot configure any AAA rules that use the local database in the system execution space.



Note

You cannot use the local database for network access authorization.

Fallback Support

The local database can act as a fallback method for several functions. This behavior is designed to help you prevent accidental lockout from the ASA.

When a user logs in, the servers in the group are accessed one at a time, starting with the first server that you specify in the configuration, until a server responds. If all servers in the group are unavailable, the ASA tries the local database if you have configured it as a fallback method (for management authentication and authorization only). If you do not have a fallback method, the ASA continues to try the AAA servers.

For users who need fallback support, we recommend that their usernames and passwords in the local database match their usernames and passwords on the AAA servers. This practice provides transparent fallback support. Because the user cannot determine whether a AAA server or the local database is providing the service, using usernames and passwords on AAA servers that are different than the usernames and passwords in the local database means that the user cannot be certain which username and password should be given.

The local database supports the following fallback functions:

- Console and enable password authentication—If the servers in the group are all unavailable, the ASA uses the local database to authenticate administrative access, which can also include enable password authentication.
- Command authorization—If the TACACS+ servers in the group are all unavailable, the local database is used to authorize commands based on privilege levels.
- VPN authentication and authorization—VPN authentication and authorization are supported to enable remote access to the ASA if AAA servers that normally support these VPN services are unavailable. When a VPN client of an administrator specifies a tunnel group configured to fallback to the local database, the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured with the necessary attributes.

How Fallback Works with Multiple Servers in a Group

If you configure multiple servers in a server group and you enable fallback to the local database for the server group, fallback occurs when no server in the group responds to the authentication request from the ASA. To illustrate, consider this scenario:

You configure an LDAP server group with two Active Directory servers, server 1 and server 2, in that order. When the remote user logs in, the ASA attempts to authenticate to server 1.

If server 1 responds with an authentication failure (such as *user not found*), the ASA does not attempt to authenticate to server 2.

If server 1 does not respond within the timeout period (or the number of authentication attempts exceeds the configured maximum), the ASA tries server 2.

If both servers in the group do not respond, and the ASA is configured to fall back to the local database, the ASA tries to authenticate to the local database.

Guidelines for the Local Database

Make sure that you prevent a lockout from the ASA when using the local database for authentication or authorization.

Related Topics[Recover from a Lockout, page 34-24](#)

Add a User Account to the Local Database

To add a user to the local database, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > Users/AAA > User Accounts**, then click **Add**.
The **Add User Account-Identity** dialog box appears.
- Step 2** Enter a username from 4 to 64 characters long.
- Step 3** Enter a password between 3 and 32 characters. Passwords are case-sensitive. The field displays only asterisks. To protect security, we recommend a password length of at least 8 characters.
-  **Note** To configure the enable password from the User Accounts pane, change the password for the enable_15 user. The enable_15 user is always present in the User Accounts pane, and represents the default username. This method of configuring the enable password is the only method available in ASDM for the system configuration. If you configured other enable level passwords at the CLI (enable password 10, for example), then those users are listed as enable_10, and so on.
- Step 4** Reenter the password.
For security purposes, only asterisks appear in the password fields.
- Step 5** Check the **User authenticated using MSCHAP** check box if you are using MSCHAP for authentication.
- Step 6** Set the management access level for a user in the **Access Restriction** area. You must first enable management authorization by clicking the **Perform authorization for exec shell access** option on the **Configuration > Device Management > Users/AAA > AAA Access > Authorization** tab.
Choose one of the following options:
- **Full Access (ASDM, Telnet, SSH and console)**—If you configure authentication for management access using the local database, then this option lets the user use ASDM, SSH, Telnet, and the console port. If you also enable authentication, then the user can access global configuration mode.
 - **Privilege Level**—Sets the privilege level for ASDM and local command authorization. The range is 0 (lowest) to 15 (highest). Specify 15 to grant unrestricted admin access. The predefined ASDM roles use 15 for Admin, 5 for Read Only, and 3 for Monitor Only (which restricts the user to the Home and Monitoring panes).
 - **CLI login prompt for SSH, Telnet and console (no ASDM access)**—If you configure authentication for management access using the local database, then this option lets the user use SSH, Telnet, and the console port. The user cannot use ASDM for configuration (if you configure HTTP authentication). ASDM monitoring is allowed. If you also configure enable authentication, then the user cannot access global configuration mode.
 - **No ASDM, SSH, Telnet, or console access**—If you configure authentication for management access using the local database, then this option disallows the user from accessing any management access method for which you configured authentication (excluding the Serial option; serial access is allowed).

Step 7 (Optional) To enable public key authentication for SSH connections to the ASA on a per-user basis, click one of the following options in the **Navigation** pane:

- **Public Key Authentication**—Paste in a Base64-encoded public key. You can generate the key using any SSH key generation software (such as `ssh-keygen`) that can generate SSH-RSA raw keys (with no certificates). When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, check the **Key is hashed** check box.

To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.

- **Public Key Using PKF**—Check the **Specify a new PKF key** check box, and paste or import a public key file (PKF) formatted key, up to 4096 bits. Use this format for keys that are too large to paste in Base64 format. For example, you can generate a 4096-bit key using `ssh-keygen`, then convert it to PKF, and import on this pane. When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, copy it from the **Public Key Authentication** pane, and paste it in that pane on the new ASA with the **Key is hashed** check box checked.

To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.

Step 8 Click **VPN Policy** to configure VPN policy attributes for this user. See the VPN configuration guide.

Step 9 Click **Apply**.

The user is added to the local database, and the changes are saved to the running configuration.



Tip

You can search for specific text in each column of the **Configuration > Device Management > Users/AAA > User Accounts** pane. Enter the specific text that you want to locate in the **Find** box, then click the **Up** or **Down** arrow. You can also use the asterisk (“*”) and question mark (“?”) as wild card characters in the text search.

Generate a Shared Key

To generate a shared key for SSH on a Linux or Macintosh system, and import it to perform the following steps:

Procedure

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
jUT4ehXIUE7FrChffBBtbD4d9FkV8A2gwZCDJBxEM26ocbZCSTx9QC//wt6E/zRcdqiJG
p4ECEdDaM+56l+yf73NUigO7wYkqcrzjmI1rZRDLCvtj8Q9qD3MqsV+PkJGSGiqZwnyI1
QbfYxXHU9wLdWxhUbA/xOjJuZ15TQMa7KLS2u+RtrpQgeTGTFfIh6O+xKh93gwTgzaZTK4
CQ1kuMrRdNRza0byLeYpTslv6Lv6F6dGtwlqrX5a+w/tV/aw9WUg/rapekKloz3tsPTDe
p866AFzU+Z7pVR1389iNuNjHQS7IUA2m0cciuCM2we/tVqMPYJl+xgKakuHDkBlMS4i8b
Wzyd+4EUMDGGZVeO+corkTLWFO1wIuieRkrUaCzjComGYZdzrQT2mXBcSKQNWlSCBpCHsk
/r5uTGnKpCNWfL7vd/sRCHyHKsxjsXR15C/5zgHmCTAaGOuIq0Rjo34+61+70PctYXebxM
Wwm19e3eH2PudZd+rj1dedfr2/IrisLEBRJWGLoR/N+xsvvVVM1Qqwlul4r99CbZF9NghY
NRxCQOY/7K77IQ==
---- END SSH2 PUBLIC KEY ----
jcrichon-mac:~.ssh john$

```

Step 3 Copy the key to your clipboard.

Step 4 In ASDM, choose **Configuration > Device Management > Users/AAA > User Accounts**, select the username and then click **Edit**. Click **Public Key Using PKF** and paste the key into the window:



Step 5 Verify the user (test) can SSH to the ASA:

```

jcrichon-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>
```

Test Local Database Authentication and Authorization

To determine whether the ASA can contact a local database and authenticate or authorize a user, perform the following steps:

Procedure

- Step 1** From the **Configuration > Device Management > Users/AAA > AAA Server Groups > AAA Server Groups** table, click the server group in which the server resides.
- Step 2** Click the server that you want to test from the **Servers in the Selected Group** table.
- Step 3** Click **Test**.
The **Test AAA Server** dialog box appears for the selected server.
- Step 4** Click the type of test that you want to perform—**Authentication** or **Authorization**.
- Step 5** Enter a username.
- Step 6** If you are testing authentication, enter the password for the username.
- Step 7** Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, ASDM displays an error message.

Monitoring the Local Database

See the following screens for monitoring the local database:

- **Monitoring > Properties > AAA Servers**

This pane shows AAA server statistics.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for the Local Database

Table 30-1 History for the Local Database

Feature Name	Platform Releases	Description
Local database configuration for AAA	7.0(1)	Describes how to configure the local database for AAA use. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > User Accounts.
Support for SSH public key authentication	9.1(2)	You can now enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF <i>Also available in 8.4(4.1); PKF key format support is only in 9.1(2).</i>



RADIUS Servers for AAA

This chapter describes how to configure RADIUS servers for AAA.

- [About RADIUS Servers for AAA, page 31-1](#)
- [Guidelines for RADIUS Servers for AAA, page 31-13](#)
- [Configure RADIUS Servers for AAA, page 31-14](#)
- [Test RADIUS Server Authentication and Authorization, page 31-18](#)
- [Monitoring RADIUS Servers for AAA, page 31-18](#)
- [History for RADIUS Servers for AAA, page 31-19](#)

About RADIUS Servers for AAA

The Cisco 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, 7.x, and 8.x.
- Microsoft

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 31-1 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 the following table 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 31-1 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
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

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
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: <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
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

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
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).

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
IPsec-Split-Tunneling-Policy	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Local LAN permitted
IPsec-Split-Tunnel-List	Y	27	String	Single	Specifies the name of the network or ACL that describes the split tunnel inclusion list.
IPsec-Tunnel-Type	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPsec-User-Group-Lock		33	Boolean	Single	0 = Disabled 1 = Enabled
IPv6-Address-Pools	Y	218	String	Single	Name of IP local pool-IPv6
IPv6-VPN-Filter	Y	219	String	Single	ACL value
L2TP-Encryption		21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Req 15= 40/128-Encr/Stateless-Req
L2TP-MPPC-Compression		38	Integer	Single	0 = Disabled 1 = Enabled
Member-Of	Y	145	String	Single	Comma-delimited string, for example: Engineering, Sales An administrative attribute that can be used in dynamic access policies. It does not set a group policy.
MS-Client-Subnet-Mask	Y	63	Boolean	Single	An IP address
NAC-Default-ACL		92	String		ACL
NAC-Enable		89	Integer	Single	0 = No 1 = Yes
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

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
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
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.

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
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"

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/ Type	Single or Multi-Valued	Description or Value
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
WebVPN-Default-Homepage	Y	76	String	Single	A URL such as http://example-example.com
WebVPN-Deny-Message	Y	116	String	Single	Valid string (up to 500 characters)
WebVPN-Download_Max-Size	Y	157	Integer	Single	0x7fffffff
WebVPN-File-Access-Enable	Y	94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Group-based-HTTP/HTTPS-Proxy-Exception-List	Y	78	String	Single	Comma-separated DNS/IP with an optional wildcard (*) (for example *.cisco.com, 192.168.1.*, wwwin.cisco.com)

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
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
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-List	Y	72	String	Single	Port forwarding list name
WebVPN-Port-Forwarding-Name	Y	79	String	Single	String name (example, "Corporate-Apps"). This text replaces the default string, "Application Access," on the clientless portal home page.
WebVPN-Post-Max-Size	Y	159	Integer	Single	0x7fffffff
WebVPN-Session-Timeout-Alert-Interval	Y	149	Integer	Single	0-30. 0 = Disabled.

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
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.
WebVPN-SVC-Gateway-DPD-Frequency	Y	109	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-Rekey-Time	Y	110	Integer	Single	4-10080 minutes, 0=Off
WebVPN-SVC-Rekey-Method	Y	111	Integer	Single	0 (Off), 1 (SSL), 2 (New Tunnel)
WebVPN-SVC-Compression	Y	112	Integer	Single	0 (Off), 1 (Deflate Compression)
WebVPN-UNIX-Group-ID (GID)	Y	222	Integer	Single	Valid UNIX group IDs
WebVPN-UNIX-User-ID (UIDs)	Y	221	Integer	Single	Valid UNIX user IDs
WebVPN-Upload-Max-Size	Y	158	Integer	Single	0x7fffffff
WebVPN-URL-Entry-Enable	Y	93	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-List	Y	71	String	Single	URL list name

Table 31-1 Supported RADIUS Authorization Attributes (continued)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
WebVPN-User-Storage	Y	160	String	Single	
WebVPN-VDI	Y	163	String	Single	List of settings

Supported IETF RADIUS Authorization Attributes

The following table lists the supported IETF RADIUS attributes.

Table 31-2 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): <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
IETF-Radius-Service-Type	Y	6	Integer	Single	Seconds. Possible Service Type values: <i>.Administrative</i>—User is allowed access to the configure prompt. <i>.NAS-Prompt</i>—User is allowed access to the exec prompt. <i>.remote-access</i>—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

Disconnect Reason Code (continued)

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

Guidelines for RADIUS Servers for AAA

This section describes the guidelines and limitations that you should check before configuring RADIUS servers for AAA.

IPv6

The AAA server must use an IPv4 address, but endpoints can use 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.

Related Topics

- [Fallback Support, page 30-4](#)
- [How Fallback Works with Multiple Servers in a Group, page 30-4](#)
- [Recover from a Lockout, page 34-24](#)

Configure RADIUS Servers for AAA

This section describes how to configure RADIUS servers for AAA.

-
- 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 [Configure RADIUS Server Groups, page 31-14](#).
- Step 3** For a server group, add a server to the group. See [Add a RADIUS Server to a Group, page 31-15](#).
- Step 4** (Optional) Specify text to display to the user during the AAA authentication challenge process. See [Add an Authentication Prompt, page 31-17](#).
-

Configure 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:

Procedure

-
- Step 1** Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**.
- Step 2** Click **Add** in the **AAA Server Groups** area.
- The **Add AAA Server Group** dialog box appears.
- Step 3** Enter a name for the group in the **Server Group** field.
- Step 4** Choose the RADIUS server type from the **Protocol** drop-down list.
- Step 5** Click **Simultaneous** or **Single** in the **Accounting Mode** field.
- In Single mode, the ASA sends accounting data to only one server.
- In Simultaneous mode, the ASA sends accounting data to all servers in the group.
- Step 6** Click **Depletion** or **Timed** in the **Reactivation Mode** field.
- In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive.
- In Timed mode, failed servers are reactivated after 30 seconds of down time.
- Step 7** If you chose the Depletion reactivation mode, enter a time interval in the **Dead Time** field.
- The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers.
- Step 8** Add the number of failed attempts allowed in the **Max Failed Attempts** field.

This option sets the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive.

- Step 9** (Optional) If you are adding a RADIUS server type, perform the following steps:
- a. Check the **Enable interim accounting update** check box if you want to enable multi-session accounting for clientless SSL VPN and AnyConnect sessions.
 - b. Check the **Enable Active Directory Agent Mode** check box to specify the shared secret between the ASA and the AD agent and indicate that a RADIUS server group includes AD agents that are not full-function RADIUS servers. Only a RADIUS server group that has been configured using this option can be associated with user identity.
 - c. Check the **Enable dynamic authorization** check box to enable ISE to send Change of Authorization (CoA) RADIUS packets. This enables policy changes made on the ISE to be enforced during the lifetime of the VPN connection.
 - d. Enter the **Dynamic Authorization Port**. This is the listening port for RADIUS CoA requests. Typically it is 1700. The valid range is 1 to 65535.
 - e. Check the **Use authorization only mode (no common password configuration required)** check box to enable authorize-only mode for the RADIUS server group. When this check box is checked, the common password configured for individual AAA servers is not required and does not need to be configured.
 - f. Click the **VPN3K Compatibility Option** down arrow to expand the list, and click one of the following options to specify whether or not a downloadable ACL received from a RADIUS packet should be merged with a Cisco AV pair ACL:
 - Do not merge
 - Place the downloadable ACL after Cisco AV-pair ACL
 - Place the downloadable ACL before Cisco AV-pair ACL
- Step 10** Click **OK**.
- The **Add AAA Server Group** dialog box closes, and the new server group is added to the **AAA Server Groups** table.
- Step 11** In the **AAA Server Groups** dialog box, click **Apply** to save the changes to the running configuration.
-

Add a RADIUS Server to a Group

To add a RADIUS server to a group, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**, and in the **AAA Server Groups** area, click the server group to which you want to add a server.
- Step 2** Click **Add** in the **Servers in the Selected Group** area (lower pane).
- The **Add AAA Server Group** dialog box appears for the server group.
- Step 3** Choose the interface name on which the authentication server resides.
- Step 4** Add either a server name or IP address for the server that you are adding to the group.

Step 5 Add a timeout value or keep the default. The timeout is the length of time, in seconds, that the ASA waits for a response from the primary server before sending the request to the backup server.

Step 6 Specify how you want the ASA to handle netmasks received in downloadable ACLs. Choose from the following options:

- **Detect automatically**—The ASA attempts to determine the type of netmask expression used. If the ASA detects a wildcard netmask expression, the ASA converts it to a standard netmask expression.



Note Because some wildcard expressions are difficult to detect clearly, this setting may misinterpret a wildcard netmask expression as a standard netmask expression.

- **Standard**—The ASA assumes downloadable ACLs received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed.
- **Wildcard**—The ASA assumes downloadable ACLs received from the RADIUS server contain only wildcard netmask expressions, and it converts them all to standard netmask expressions when the ACLs are downloaded.

Step 7 Specify a case-sensitive password that is common among users who access this RADIUS authorization server through this ASA. Be sure to provide this information to your RADIUS server administrator.



Note For an authentication RADIUS server (rather than authorization), do not configure a common password.

If you leave this field blank, the username is the password for accessing this RADIUS authorization server.

Never use a RADIUS authorization server for authentication. Common passwords or usernames as passwords are less secure than assigning unique user passwords.

Although the password is required by the RADIUS protocol and the RADIUS server, users do not need to know it.

Step 8 If you use double authentication and enable password management in the tunnel group, then the primary and secondary authentication requests include MS-CHAPv2 request attributes. If a RADIUS server does not support MS-CHAPv2, then you can configure that server to send a non-MS-CHAPv2 authentication request by unchecking this check box.

Step 9 Specify the length of time, from 1 to 10 seconds, that the ASA waits between attempts to contact the server.



Note The interval between subsequent retries will be always 50ms or 100ms, regardless of the retry-interval settings you have entered. This is the intended behavior.

Step 10 Click **Simultaneous** or **Single**.

In Single mode, the ASA sends accounting data to only one server.

In Simultaneous mode, the ASA sends accounting data to all servers in the group.

Step 11 Specify the server port to be used for accounting of users. The default port is 1646.

Step 12 Specify the server port to be used for authentication of users. The default port is 1645.

- Step 13** Specify the shared secret key used to authenticate the RADIUS server to the ASA. The server secret that you configure should match the one configured on the RADIUS server. If you do not know the server secret, ask the RADIUS server administrator. The maximum field length is 64 characters.
- Step 14** Click **OK**.
- The **Add AAA Server Group** dialog box closes, and the AAA server is added to the AAA server group.
- Step 15** In the **AAA Server Groups** pane, click **Apply** to save the changes to the running configuration.

Add an Authentication Prompt

You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from RADIUS servers. This text is primarily for cosmetic purposes and appears above the username and password prompts that users see when they log in. If you do not specify an authentication prompt, users see the following when authenticating with a RADIUS server:

Connection Type	Default Prompt
FTP	FTP authentication
HTTP	HTTP authentication
Telnet	None

To add an authentication prompt, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > Users/AAA > Authentication Prompt**.
- Step 2** Enter text in the **Prompt** field to add as a message to appear above the username and password prompts that users see when they log in.

The following table shows the allowed character limits for authentication prompts:

Application	Character Limit
Microsoft Internet Explorer	37
Telnet	235
FTP	235

- Step 3** Add messages in the **User accepted message** and **User rejected message** fields.
- If the user authentication occurs from Telnet, you can use the **User accepted message** and **User rejected message** options to display different status prompts to indicate that the authentication attempt is either accepted or rejected by the RADIUS server.
- If the RADIUS server authenticates the user, the ASA displays the **User accepted message** text, if specified, to the user; otherwise, the ASA displays the **User rejected message** text, if specified. Authentication of HTTP and FTP sessions displays only the challenge text at the prompt. The **User accepted message** and **User rejected message** text are not displayed.

Step 4 Click **Apply** to save the changes to the running configuration.

Test RADIUS Server Authentication and Authorization

To determine whether the ASA can contact a RADIUS server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**.

Step 2 Click the server group in which the server resides in the **AAA Server Groups** table.

Step 3 Click the server that you want to test in the **Servers in the Selected Group** table.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 Enter the password for the username if you are testing authentication.

Step 8 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring RADIUS Servers for AAA

See the following screens for monitoring the status of RADIUS servers for AAA:

- **Monitoring > Properties > AAA Servers**

This pane shows the RADIUS server running configuration.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for RADIUS Servers for AAA

Table 31-3 History for RADIUS Servers for AAA

Feature Name	Platform Releases	Description
RADIUS Servers for AAA	7.0(1)	Describes how to configure RADIUS servers for AAA. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > Authentication Prompt.
Key vendor-specific attributes (VSAs) sent in RADIUS access request and accounting request packets from the ASA	8.4(3)	Four New VSAs—Tunnel Group Name (146) and Client Type (150) are sent in RADIUS access request packets from the ASA. Session Type (151) and Session Subtype (152) are sent in RADIUS accounting request packets from the ASA. All four attributes are sent for all accounting request packet types: Start, Interim-Update, and Stop. The RADIUS server (for example, ACS and ISE) can then enforce authorization and policy attributes or use them for accounting and billing purposes.



TACACS+ Servers for AAA

This chapter describes how to configure TACACS+ servers used in AAA.

- [About TACACS+ Servers for AAA, page 32-1](#)
- [Guidelines for TACACS+ Servers for AAA, page 32-2](#)
- [Configure TACACS+ Servers, page 32-3](#)
- [Test TACACS+ Server Authentication and Authorization, page 32-6](#)
- [Monitoring TACACS+ Servers for AAA, page 32-6](#)
- [History for TACACS+ Servers for AAA, page 32-6](#)

About TACACS+ Servers for AAA

The ASA supports TACACS+ server authentication with the following protocols: ASCII, PAP, CHAP, and MS-CHAPv1.

TACACS+ Attributes

The Cisco 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.

The following table lists supported TACACS+ authorization response attributes for cut-through-proxy connections.

Table 32-1 **Supported TACACS+ Authorization Response Attributes**

Attribute	Description
acl	Identifies a locally configured ACL to be applied to the connection.

Table 32-1 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.

The following table lists supported TACACS+ accounting attributes.

Table 32-2 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_addr	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.

Guidelines for TACACS+ Servers for AAA

This section describes the guidelines and limitation that you should check before configuring TACACS+ servers for AAA.

IPv6

The AAA server must use an IPv4 address, but endpoints can use 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.

Related Topics

- [Fallback Support, page 30-4](#)
- [How Fallback Works with Multiple Servers in a Group, page 30-4](#)
- [Recover from a Lockout, page 34-24](#)

Configure TACACS+ Servers

This section describes how to configure TACACS+ servers.

-
Step 1	Add a TACACS+ server group. See Configure TACACS+ Server Groups, page 32-3 .
Step 2	For a server group, add a server to the group. See Add a TACACS+ Server to a Group, page 32-4 .
Step 3	(Optional) Specify text to display to the user during the AAA authentication challenge process. See Add an Authentication Prompt, page 32-5 .
-

Configure 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:

Procedure

-
Step 1	Choose Configuration > Device Management > Users/AAA > AAA Server Groups .
Step 2	Click Add in the AAA Server Groups area. The Add AAA Server Group dialog box appears.
Step 3	Enter a name for the group in the Server Group field.
Step 4	Choose the TACACS+ server type from the Protocol drop-down list:
Step 5	Click Simultaneous or Single in the Accounting Mode field. In Single mode, the ASA sends accounting data to only one server. In Simultaneous mode, the ASA sends accounting data to all servers in the group.
Step 6	Click Depletion or Timed in the Reactivation Mode field. In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive. In Timed mode, failed servers are reactivated after 30 seconds of down time.
Step 7	If you chose the Depletion reactivation mode, enter a time interval in the Dead Time field.

The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers.

Step 8 Add the maximum number of failed attempts allowed.

This option sets the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive.

Step 9 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the new server group is added to the **AAA Server Groups** table.

Step 10 Click **Apply** to save the changes to the running configuration.

Add a TACACS+ Server to a Group

To add a TACACS+ server to a group, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**.

Step 2 Click the server group to which you want to add a server.

Step 3 Click **Add** in the **Servers in the Selected Group** area.

The **Add AAA Server Group** dialog box appears for the server group.

Step 4 Choose the interface name on which the authentication server resides.

Step 5 Add either a server name or IP address for the server that you are adding to the group.

Step 6 Either add a timeout value or keep the default. The timeout is the duration of time, in seconds, that the ASA waits for a response from the primary server before sending the request to the backup server.

Step 7 Specify the server port. The server port is either port number 139, or the TCP port number used by the ASA to communicate with the TACACS+ server.

Step 8 Specify the server secret key. The shared secret key used to authenticate the TACACS+ server to the ASA. The server secret that you configure here should match the one that is configured on the TACACS+ server. If you do not know the server secret, ask the TACACS+ server administrator. The maximum field length is 64 characters.

Step 9 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the AAA server is added to the AAA server group.

Step 10 Click **Apply** to save the changes to the running configuration.

Add an Authentication Prompt

You can specify text to display to the user during the AAA authentication challenge process. You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from TACACS+ servers. This text is primarily for cosmetic purposes and appears above the username and password prompts that users see when they log in.

If you do not specify an authentication prompt, users see the following when authenticating with a TACACS+ server:

Connection Type	Default Prompt
FTP	FTP authentication
HTTP	HTTP authentication
Telnet	None

To add an authentication prompt, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > Authentication Prompt**.

Step 2 Add text to appear above the username and password prompts that users see when they log in.

The following table shows the allowed character limits for authentication prompts:

Application	Character Limit for Authentication Prompt
Microsoft Internet Explorer	37
Telnet	235
FTP	235

Step 3 Add messages in the **User accepted message** and **User rejected message** fields.

If the user authentication occurs from Telnet, you can use the **User accepted message** and **User rejected message** options to display different status prompts to indicate that the authentication attempt is accepted or rejected by the AAA server.

If the AAA server authenticates the user, the ASA displays the **User accepted message** text, if specified, to the user; otherwise, the ASA displays the **User rejected message** text, if specified. Authentication of HTTP and FTP sessions displays only the challenge text at the prompt. The User accepted message and User rejected message text are not displayed.

Step 4 Click **Apply** to save the changes to the running configuration.

Test TACACS+ Server Authentication and Authorization

To determine whether the ASA can contact a TACACS+ server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**.

Step 2 Click the server group in which the server resides.

Step 3 Click the server that you want to test.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 If you are testing authentication, enter the password for the username.

Step 8 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring TACACS+ Servers for AAA

See the following screens for monitoring TACACS+ servers for AAA:

- **Monitoring > Properties > AAA Servers**

This pane shows the configured TACACS+ server statistics.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for TACACS+ Servers for AAA

Table 32-3 History for TACACS+ Servers for AAA

Feature Name	Platform Releases	Description
TACACS+ Servers	7.0(1)	Describes how to configure TACACS+ servers for AAA. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > Authentication Prompt.



LDAP Servers for AAA

This chapter describes how to configure LDAP servers used in AAA.

- [About LDAP and the ASA, page 33-1](#)
- [Guidelines for LDAP Servers for AAA, page 33-4](#)
- [Configure LDAP Servers for AAA, page 33-5](#)
- [Test LDAP Server Authentication and Authorization, page 33-8](#)
- [Monitoring LDAP Servers for AAA, page 33-9](#)
- [History for LDAP Servers for AAA, page 33-9](#)

About LDAP and the ASA

The Cisco 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.

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 LDAP, see RFCs 1777, 2251, and 2849.

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 the following figure 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 33-1 A Multi-Level LDAP Hierarchy



330368

Search 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.

The figure shows a sample LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. The following table 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 33-1 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

Bind 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.

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.

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.

Guidelines for LDAP Servers for AAA

This section includes the guidelines and limitations that you should check before configuring LDAP servers for AAA.

IPv6

The AAA server must use an IPv4 address, but endpoints can use IPv6.

Additional 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.
- 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.

Configure LDAP Servers for AAA

This section describes how to configure LDAP servers for AAA.

-
Step 1	Configure LDAP attribute maps. See Configure LDAP Attribute Maps, page 33-5 .
Step 2	Add an LDAP server group. See Configure LDAP Server Groups, page 33-6 .
Step 3	Add a server to the group, then configure server parameters. See Add an LDAP Server to a Server Group, page 33-7 .
-

Configure LDAP Attribute Maps

To configure LDAP attribute maps, perform the following steps:

Procedure

-
Step 1	Choose Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map (for local users), or Configuration > Device Management > Users/AAA > LDAP Attribute Map (for all other users), then click Add . The Add LDAP Attribute Map dialog box appears with the Mapping of Attribute Name tab active.
Step 2	Create a name for this attribute map.
Step 3	Add the name of one of the LDAP attributes to be mapped.
Step 4	Choose a Cisco attribute.
Step 5	Click Add .
Step 6	To map more attributes, repeat Steps 1 through 5.

- Step 7** Click the **Mapping of Attribute Value** tab to map the value of any of the LDAP attributes to a new value in the mapped Cisco attribute.
 - Step 8** Click **Add**.
The **Add Mapping of Attribute Value** dialog box appears.
 - Step 9** Enter the value for this LDAP attribute that you expect to be returned from the LDAP server.
 - Step 10** Enter the value you want to use in the Cisco attribute when this LDAP attribute contains the previous LDAP attribute value.
 - Step 11** Click **Add**.
 - Step 12** To map more attribute values, repeat Steps 8 through 11.
 - Step 13** Click **OK** twice to close each dialog box.
 - Step 14** Click **Apply** to save the settings to the running configuration.
-

Configure LDAP Server Groups

This section describes how to configure LDAP server groups.

Before You Begin

You must add an attribute map before you may add an LDAP server to an LDAP server group.

To create and configure an LDAP server group, then add an LDAP server to that group, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**, or **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups** for VPN users.
 - Step 2** Click **Add**.
The **Add AAA Server Group** dialog box appears.
 - Step 3** Enter a name for the AAA server group.
 - Step 4** Choose the LDAP server type from the **Protocol** drop-down list.
 - Step 5** Click the radio button for the reactivation mode you want to use (**Depletion** or **Timed**).
In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive.
In Timed mode, failed servers are reactivated after 30 seconds of down time.
 - a.** If you chose the Depletion reactivation mode, enter a time interval in the **Dead Time** field.
The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers.
 - Step 6** Add the maximum number of failed attempts to connect to the server to allow.
This option sets the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive.
 - Step 7** Click **OK**.

The **Add AAA Server Group** dialog box closes, and the new server group is added to the AAA server group.

- Step 8** Click **Apply** to save the changes to the running configuration.

Add an LDAP Server to a Server Group

To add an LDAP server to a server group, perform the following steps:

Procedure

-
- Step 1** Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**, or **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups** for VPN users.
- Step 2** Select the server group to which you want to add a server, then click **Add**.
The **Add AAA Server** dialog box appears for the selected server group.
- Step 3** Choose the name of the interface that connects to the LDAP server.
- Step 4** Add either the server name or IP address of the LDAP server.
- Step 5** Add a timeout value or keep the default. The timeout is the duration of time, in seconds, that the ASA waits for a response from the primary server before sending the request to the backup server.
- Step 6** In the **LDAP Parameters for authentication/authorization** area, configure the following settings:
- **Enable LDAP over SSL** (also called secure LDAP or LDAP-S)—Check this check box to use SSL to secure communications between the ASA and the LDAP server.



Note If you do not configure the SASL protocol, we strongly recommend that you secure LDAP communications with SSL.

- **Server Port**—Enter TCP port number 389, the port which the ASA uses to access the LDAP server for simple (non-secure) authentication, or TCP port 636 for secure authentication (LDAP-S). All LDAP servers support authentication and authorization. Only Microsoft AD and Sun LDAP servers additionally provide a VPN remote access password management capability, which requires LDAP-S.
- **Server Type**—Specify the LDAP server type from the drop-down list. The available options include the following:
 - **Detect Automatically/Use Generic Type**
 - **Microsoft**
 - **Novell**
 - **OpenLDAP**
 - **Sun, now part of Oracle Directory Server Enterprise Edition**
- **Base DN**—Enter the base distinguished name (DN) or location in the LDAP hierarchy where the server should begin searching when it receives an LDAP request (for example, OU=people, dc=cisco, dc=com).
- **Scope**—Specify the extent of the search that the server should perform in the LDAP hierarchy when it receives an authorization request from the drop-down list. The following options are available:

- **One Level**—Searches only one level beneath the Base DN. This option is quicker.
- **All Levels**—Searches all levels beneath the Base DN (that is, searches the entire subtree hierarchy). This option takes more time.
- **Naming Attribute(s)**—Enter the relative distinguished name attribute(s) that uniquely identify an entry on the LDAP server. Common naming attributes are Common Name (CN), sAMAccountName, userPrincipalName, and User ID (uid).
- **Login DN and Login Password**—The ASA uses the login DN and login password to establish trust (bind) with an LDAP server. Specify the login password, which is the password for the login DN user account.
- **LDAP Attribute Map**—Select one of the attribute maps that you created for this LDAP server to use. These attribute maps map LDAP attribute names to Cisco attribute names and values.
- **SASL MD5 authentication**—This enables the MD5 mechanism of the SASL to authenticate communications between the ASA and the LDAP server.
- **SASL Kerberos authentication**—Enables the Kerberos mechanism of the SASL to secure authentication communications between the ASA and the LDAP server. You must have defined a Kerberos server in order to enable this option.
- **LDAP Parameters for Group Search**—The fields in this area configure how the ASA requests AD groups.
 - **Group Base DN**—Specifies the location in the LDAP hierarchy to begin searching for the AD groups (that is, the list of memberOf enumerations). If this field is not configured, the ASA uses the base DN for AD group retrieval. ASDM uses the list of retrieved AD groups to define AAA selection criteria for dynamic access policies. For more information, see the **show ad-groups** command.
 - **Group Search Timeout**—Specify the maximum time to wait for a response from an AD server that was queried for available groups.

Step 7 Click **OK**.

The **Add AAA Server** dialog box closes, and the AAA server is added to the AAA server group.

Step 8 Click **Apply** to save the changes to the running configuration.

Test LDAP Server Authentication and Authorization

To determine whether the ASA can contact an LDAP server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Server Groups**.

Step 2 Select the server group in which the server resides.

Step 3 Select the server that you want to test.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 If you are testing authentication, enter the password for the username.

Step 8 Click **OK**.

The ASA sends either an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring LDAP Servers for AAA

See the following screens for monitoring LDAP servers for AAA:

- **Monitoring > Properties > AAA Servers**

This pane shows the configured AAA server statistics.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for LDAP Servers for AAA

Table 33-2 History for AAA Servers

Feature Name	Platform Releases	Description
LDAP Servers for AAA	7.0(1)	LDAP Servers describe support for AAA and how to configure LDAP servers. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map.



PART 7

System Administration



Management Access

This chapter describes how to access the Cisco ASA for system management through Telnet, SSH, and HTTPS (using ASDM), how to authenticate and authorize users, and how to create login banners.

- [Guidelines for Management Access, page 34-1](#)
- [ASDM Authorization and Username from Certificate Extraction, page 34-3](#)
- [Configure ASA Access for ASDM, Telnet, or SSH, page 34-4](#)
- [Configure AAA for System Administrators, page 34-9](#)
- [Monitoring Device Access, page 34-25](#)
- [History for Management Access, page 34-26](#)

Guidelines for Management Access

This section describes the guidelines and limitation that you should check before configuring management access.

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.

VPN Guidelines



Note

For the configurations that follow, 192.168.10.0/24 is the VPN pool for AnyConnect or IPsec VPN clients. Each configuration allows VPN client users to connect to ASDM or SSH to the ASA using the management interface IP address.

- To allow only VPN client users access to ASDM or HTTP (and deny access to all other users), enter the following commands:

```
ciscoasa(config)# http server enable  
ciscoasa(config)# http 192.168.10.0 255.255.255.0 management_interface
```

- To allow only VPN client users access to the ASA using SSH (and deny access to all other users), enter the following command:

```
ciscoasa(config)# ssh 192.168.10.0 255.255.255.0 management_interface
```

- You can define only one management access interface.

Additional Guidelines

- 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.
- 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.
- 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 by choosing **Configuration > Device Management > Users/AAA > AAA Access > Authentication**; then define a local user by choosing **Configuration > Device Management > Users/AAA > User Accounts**. If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.
- (9.1(2) and later) The default Telnet login password was removed; you must manually set the password before using Telnet.
- To gain access to the ASA CLI using Telnet, enter the login password. You must manually set the password before using Telnet.
- If you configure Telnet authentication, then enter the username and password defined by the AAA server or local database.
- 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.

- If you cannot make a Telnet or SSH connection to the ASA interface, make sure that you enabled Telnet or SSH to the ASA according to the instructions in this chapter.
- 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.

- After a banner has been added, Telnet or SSH sessions to the 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).
- See RFC 2196 for guidelines about banner messages.

Related Topics

- [Set the Hostname, Domain Name, and the Enable and Telnet Passwords, page 18-1](#)
- [Add a User Account to the Local Database, page 30-5](#)

ASDM Authorization and Username from Certificate Extraction

Until this release, the ASA supported certificate-only authentication, AAA authentication using a username and password entered on the ASDM client, and both certificate and AAA authentication when the ASDM connection is established.

In this release, with the addition of the ASDM Authorization and Username from Certificate Extraction feature, we now support:

- Certificate authentication with User authorization using username from certificate
- AAA authentication with User authorization using username provided by the user
- Certificate authentication and AAA authentication with User authorization using username provided by the user
- Certificate authentication and AAA authentication with prefilled username using username from certificate
- Certificate authentication and AAA authentication with prefilled username using username from certificate with User authorization using the same username

Set the Rules for Extracting a Username from the Certificate

An administrator can configure the rules for extracting a username from a certificate and choose to specify if the username that is extracted from the certificate is used for authentication and pre-filled in the form presented to the user (pre-fill username).

-



Note

Step 1 Browse to **Configuration > Device Management > Management Access > HTTP Certificate Rule**.



Note

This configuration is not available in transparent firewall and multi-context modes.

Step 2 Click **Specify the certificate fields to be used as the username** to enter primary and second fields.

Step 3 Choose from the Primary Field and Secondary Field drop-down values to specify which attributes and additional attributes to use to derive the username.

- C—Country: the two-letter country abbreviation which conforms to the ISO 3166 country abbreviations.
- CN—Common Name: the name of a person, system, or other entity. Not available as a secondary attribute.
- DNQ—Domain Name Qualifier.
- EA—Email address.
- GENQ—Generational qualifier.
- GN—Given name.
- I—Initials.
- L—Locality: the city or town where the organization is located.
- N—Name.
- O—Organization: the name of the company, institution, agency, association, or other entity.
- OU—Organizational Unit: the subgroup within the organization (O).
- SER—Serial number.
- SN—Surname.
- SP—State/Province: the state or province where the organization is located.
- T—Title.
- UID—User Identifier.
- UPN—User Principal Name.

Step 4 (Optional for Primary Attribute Only) Choose **Use the entire DN as the username**.

Step 5 (Optional) Choose **Use script to select username** if you want to use the LUA script generated by ASDM.

Step 6 Check the **Prefill Username** checkbox to enable the use of this name for authentication. When enabled, this username, along with the password entered by the user, are used for authentication. If the username is different from the one initially entered by the user, a new pop-up appears with the username pre-filled, requiring the user to only enter a password for authentication.

Configure ASA Access for ASDM, Telnet, or SSH

This section describes how to configure ASA access for ASDM, Telnet, or SSH.

Before You Begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

To identify the client IP addresses allowed to connect to the ASA using Telnet, SSH, or ASDM, perform the following steps.

Procedure

-
- Step 1** Choose **Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH**, and click **Add**.
- The **Add Device Access Configuration** dialog box appears.
- Step 2** Choose the type of session from the three options listed: **ASDM/HTTPS**, **Telnet**, or **SSH**.
- Step 3** Choose the management interface and set the host IP addresses allowed, and click **OK**.
- Step 4** Make sure the **Enable HTTP Server** check box is checked. This is enabled by default. Set other HTTP server options as needed.
- Step 5** (Optional) Configure Telnet settings. The default timeout value is 5 minutes.
- Step 6** (Optional) Configure SSH settings. For **DH Key Exchange**, click the applicable radio button to choose Diffie-Hellman (DH) Key Exchange Group 1 or Group 14. Both the DH Group 1 and Group 14 key-exchange methods for key exchange are supported on the ASA. If no DH group key-exchange method is specified, the DH group 1 key-exchange method is used. For more information about using DH key-exchange methods, see RFC 4253.
- Step 7** Click **Apply**.
- Step 8** (Required for Telnet) Set a login password before you can connect with Telnet; there is no default password.
- Choose **Configuration > Device Setup > Device Name/Password**.
 - Check the **Change the password to access the console of the security appliance** check box in the **Telnet Password** area.
 - Enter the old password (leave this field blank for a new ASA), new password, then confirm the new password.
 - Click **Apply**.
- Step 9** (Required for SSH) Configure SSH user authentication.
- Choose **Configuration > Device Management > Users/AAA > AAA Access > Authentication**.
 - Check the **SSH** check box.
 - Choose the **LOCAL** database from the **Server Group** drop-down list. You can alternatively configure authentication using a AAA server.
 - Click **Apply**.
 - Add a local user. Choose **Configuration > Device Management > Users/AAA > User Accounts**, then click **Add**.
- The **Add User Account-Identity** dialog box appears.
- Enter a username and password, then confirm the password.
 - Click **OK**, then click **Apply**.
-

Configure HTTP Redirect

You use HTTPS to connect to the ASA using ASDM. For your convenience, you can redirect HTTP connections to the management interface to HTTPS. For example, by redirecting HTTP, users can enter either `http://10.1.1.4/admin/` or `https://10.1.8.4/admin/` and still arrive at the ASDM launch page at the HTTPS address.



Tip

The access rules on the management interface must allow both HTTP and HTTPS connections; these protocols normally use port 80 and 443, respectively.

To enable redirection for each interface that you support for ASDM access, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > HTTP Redirect**.
The table shows the currently-configured interfaces and whether redirection is enabled on a interface.
- Step 2** Select the interface that you use for ASDM, and click **Edit**.
- Step 3** Configure the following options in the **Edit HTTP/HTTPS Settings** dialog box:
 - **Redirect HTTP to HTTPS**—Redirects HTTP requests to HTTPS.
 - **HTTP Port**—Identifies the port from which the interface redirects HTTP connections. The default is 80.
- Step 4** Click **OK**.

Configure 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.

To configure a login banner, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > Management Access > Command Line (CLI) > Banner**.
- Step 2** Add your banner text to the field for the type of banner that you are creating for the CLI:
 - The session (exec) banner appears when a user accesses privileged EXEC mode at the CLI.
 - The login banner appears when a user logs in to the CLI.
 - The message-of-the-day (motd) banner appears when a user first connects to the CLI.
 - The ASDM banner appears when a user connects to ASDM, after user authentication. The user is given two options for dismissing the banner:
 - **Continue**—Dismiss the banner and complete login.
 - **Disconnect**—Dismiss the banner and terminate the connection.

- Only ASCII characters are allowed, including a new line (Enter), which counts as two characters.
- Do not use tabs in the banner, because they are not preserved in the CLI version.
- There is no length limit for banners other than those for RAM and flash memory.
- You can dynamically add the hostname or domain name of the ASA by including the strings **\$(hostname)** and **\$(domain)**.
- If you configure a banner in the system configuration, you can use that banner text within a context by using the **\$(system)** string in the context configuration.

Step 3 Click **Apply**.

The new banner is saved to the running configuration.

Customize a CLI Prompt

The **CLI Prompt** pane lets you customize the prompt used during CLI sessions. By default, the prompt shows the hostname of the ASA. In multiple context mode, the prompt also displays the context name. You can display the following items in the CLI prompt:

cluster-unit	(Single and multiple mode) Displays the cluster unit name. Each unit in a cluster can have a unique name.
context	(Multiple mode only) Displays the name of the current context.
domain	Displays the domain name.
hostname	Displays the hostname.
priority	Displays the failover priority as pri (primary) or sec (secondary).
state	Displays the traffic-passing state of the unit. The following values appear for the state: • act—Failover is enabled, and the unit is actively passing traffic. • stby—Failover is enabled, and the unit is not passing traffic and is in a standby, failed, or another inactive state. • actNoFailover—Failover is not enabled, and the unit is actively passing traffic. • stbyNoFailover—Failover is not enabled, and the unit is not passing traffic. This condition might occur when there is an interface failure above the threshold on the standby unit. Shows the role (master or slave) of a unit in a cluster. For example, in the prompt ciscoasa/cl2/slave, the hostname is ciscoasa, the unit name is cl2, and the state name is slave.

To customize the CLI prompt, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > Management Access > Command Line (CLI) > CLI Prompt**.

- Step 2** Do any of the following to customize the prompt:
- Click the attribute in the **Available Prompts** list, then click **Add**. You can add multiple attributes to the prompt. The attribute is moved from the **Available Prompts** list to the **Selected Prompts** list.
 - Click the attribute in the **Selected Prompts** list, then click **Delete**. The attribute is moved from the **Selected Prompts** list to the **Available Prompts** list.
 - Click the attribute in the **Selected Prompts** list and click **Move Up** or **Move Down** to change the order in which the attributes appear.
- The prompt is changed and appears in the **CLI Prompt Preview** field.
- Step 3** Click **Apply**.
- The new prompt is saved to the running configuration.
-

Change 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:

Procedure

- Step 1** Choose **Configuration > Device Management > Management Access > Command Line (CLI) > Console Timeout**.
- Step 2** Define a new timeout value in minutes. To specify an unlimited amount of time, enter **0**. The default value is 0.
- Step 3** Click **Apply**.
- The timeout value change is saved to the running configuration.
-

Configure 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.

To configure a management interface, perform the following steps.

Procedure

- Step 1** Choose **Configuration > Device Management > Management Access > Management Interface**.

Step 2 Choose the interface with the highest security (the inside interface) from the **Management Access Interface** drop-down list.

Step 3 Click **Apply**.

The management interface is assigned, and the change is saved to the running configuration.

Configure AAA for System Administrators

This section describes how to enable authentication and command authorization for system administrators.

CLI Access with and without Authentication

How you log into the ASA depends on whether or not you enable authentication:

- **No Authentication**—If you do not enable any authentication for Telnet, you do not enter a username; you enter the login password. (SSH is not available without authentication). You access user EXEC mode.
- **Authentication**—If you enable Telnet or SSH authentication according to this section, you enter the username and password as defined on the AAA server or local user database. You access user EXEC mode.

To enter privileged EXEC mode after logging in, enter the **enable** command. How **enable** works depends on whether you enable authentication:

- **No Authentication**—If you do not configure enable authentication, enter the system enable password when you enter the **enable** command. However, if you do not use enable authentication, after you enter the **enable** command, you are no longer logged in as a particular user. To maintain your username, use enable authentication.
- **Authentication**—If you configure enable authentication, the ASA prompts you for your username and password again. This feature is particularly useful when you perform command authorization, in which usernames are important in determining the commands that a user can enter.

For enable authentication using the local database, you can use the **login** command instead of the **enable** command. **login** maintains the username but requires no configuration to turn on authentication.

ASDM Access with and without Authentication

By default, you can log into ASDM with a blank username and the enable password. Note that if you enter a username and password at the login screen (instead of leaving the username blank), ASDM checks the local database for a match.

If you configure HTTP authentication, you can no longer use ASDM with a blank username and the enable password.

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.

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.

- 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.

Preserve 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.

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 cannot enter configuration mode.

Configure Authentication for CLI, ASDM, and enable command Access

This section describes how to configure authentication for CLI, ASDM, and enable command access.

Before You Begin

- Configure Telnet, SSH, or HTTP access.
- You must configure SSH authentication to obtain SSH access; there is no default username.

To configure authentication for CLI, ASDM, and enable command access, perform the following steps:

Procedure

-
- Step 1** To authenticate users who use the **enable** command, choose **Configuration > Device Management > Users/AAA > AAA Access > Authentication**, then configure the following settings:
- a. Check the **Enable** check box.
 - b. Choose a server group name or the LOCAL database.
 - c. (Optional) If you chose a AAA server, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Check the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server, because the ASA prompt does not give any indication of which method is being used.

- Step 2** To authenticate users who access the CLI or ASDM, choose **Configuration > Device Management > Users/AAA > AAA Access > Authentication**, then configure the following settings:
- Check one or more of the following check boxes:
 - HTTP/ASDM**—Authenticates the ASDM client that accesses the ASA using HTTPS. HTTP management authentication does not support the SDI protocol for a AAA server group.
 - Serial**—Authenticates users who access the ASA using the console port. For the ASASM, this parameter affects the virtual console accessed from the switch using the **service-module session** command.
 - SSH**—Authenticates users who access the ASA using SSH.
 - Telnet**—Authenticates users who access the ASA using Telnet. For the ASASM, this parameter also affects the session from the switch using the **session** command.
 - For each service that you checked, choose a server group name or the LOCAL database.
 - (Optional) If you chose a AAA server, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Check the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server because the ASA prompt does not give any indication of which method is being used.
- Step 3** Click **Apply**.
- Step 4** To configure the rules for extracting username from certificate, choose **Configuration > Device Management > Management Access > HTTP Certificate Rule**, and then do one of the following:
- Click the **Specify the Certificate Fields to be used** radio button and select a value from the **Primary Field** and the **Secondary Field** drop-down lists.
 - Click the **Use the entire DN as the username** radio button.
 - Click the **Use script to select username** and click **Add** to add the script content.
-  **Note** Check the **Prefill Username** check box to enable the username to be extracted from the certificate for authentication.
- Step 5** Click **Apply**.

Limit User CLI and ASDM Access with Management Authorization

The ASA enables you to distinguish between administrative and remote-access users when they authenticate using RADIUS, LDAP, TACACS+, or the local user database. User role differentiation can prevent remote access VPN and network access users from establishing an administrative connection to the ASA.



Note

Serial access is not included in management authorization, so if you enable the **Authentication > Serial** option, then any user who authenticates can access the console port.

Procedure

Step 1 Choose one of the following options:



Note



Note

- To enable management authorization for Telnet and SSH sessions, choose **Configuration > Device Management > Users/AAA > AAA Access > Authorization**, and check the **Enable** check box in the **Enable Authorization for ASA Command Access Area**.
- To enable management authorization for HTTP sessions, choose **Configuration > Device Management > Users/AAA > AAA Access > Authorization**, and check the **HTTP** check box in the **Enable Authorization for ASA Command Access Area**.

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.

When the **authentication-server** option is configured, the same server is used for both authentication and authorization. You can set separate configuration of authorization for HTTP or set authorization and authentication to work independently or together.

Step 2 Click **Configure Command Privileges** to assign privilege levels of commands to individuals or groups. Authorization is disabled by default, so you must click **Set ASDM Defined User Roles** to enable.

Step 3 In the **Perform authorization for exec shell access** area, check **Enable** to perform authorization and select either the **Remote** or **Local** radio buttons to specify the server to be used for authorization of exec shell access.

Step 4 To enable management authorization, check the **Allow privileged users to enter into EXEC mode on login** check box.

The **auto-enable** option allows users with sufficient privileges from the login authentication server to be placed directly in privileged EXEC mode. Otherwise, users are placed in user EXEC mode. These privileges are determined by the Service-Type and Privilege-Level attributes that are required to enter each EXEC mode. To enter privileged EXEC mode, users must have a Service-Type attribute of Administrative and a Privilege Level attribute of greater than 1 assigned to them.

This option is not supported in the system context. However, if you configure Telnet or serial authentication in the admin context, then authentication also applies to sessions from the switch to the ASASM.

There is no effect if you enter the **aaa authorization exec** command alone.

The **auto-enable** option is not included when you use serial authentication in management authorization.

The **aaa authentication http** command is not affected by the **auto-enable** option.

Before you configure the **auto-enable** option, we recommend that you configure both protocol login and enable authentication, and that all authentication requests go to the same AAA server group, as shown in the following example:

```
ciscoasa (config)# aaa authentication ssh console RADIUS
ciscoasa (config)# aaa authentication enable console RADIUS
ciscoasa (config)# aaa authorization exec authentication-server auto-enable
```

We do *not* recommend that you use other types of configurations.

Step 5 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.

The RADIUS IETF **service-type** attribute, when sent in an access-accept message as the result of a RADIUS authentication and authorization request, is used to designate which type of service is granted to the authenticated user:

- Service-Type 6 (Administrative)—Allows full access to any services specified by the **Authentication** tab options.
- Service-Type 7 (NAS prompt)—Allows access to the CLI when you configure the **Telnet** or **SSH** authentication options, but denies ASDM configuration access if you configure the **HTTP** option. ASDM monitoring access is allowed. If you configure **enable** authentication with the **Enable** option, the user cannot access privileged EXEC mode using the **enable** command. The Framed (2) and Login (1) service types are treated the same way.
- Service-Type 5 (Outbound)—Denies management access. The user cannot use any services specified by the **Authentication** tab options (excluding the **Serial** option; serial access is allowed). Remote access (IPsec and SSL) users can still authenticate and terminate their remote access sessions. All other service types (Voice, FAX, and so on) are treated the same way.

The RADIUS Cisco VSA **privilege-level** attribute (Vendor ID 3076, sub-ID 220), when sent in an access-accept message, is used to designate the level of privilege for the user.

When an authenticated user tries administrative access to the ASA through ASDM, SSH, or Telnet, but does not have the appropriate privilege level to do so, the ASA generates syslog message 113021. This message informs the user that the attempted login failed because of inappropriate administrative privileges.

- TACACS+ users

Authorization is requested with “service=shell,” and the server responds with PASS or FAIL.

- PASS, privilege level 1—Allows full access to any services specified by the **Authentication** tab options.
- PASS, privilege level 2 and higher—Allows access to the CLI when you configure the **Telnet** or **SSH** authentication options, but denies ASDM configuration access if you configure the **HTTP** option. ASDM monitoring access is allowed. If you configure **enable** authentication with the **Enable** option, the user cannot access privileged EXEC mode using the **enable** command. You are not allowed to access privileged EXEC mode using the **enable** command if your enable privilege level is set to 14 or less.
- FAIL—Denies management access. You cannot use any services specified by the **Authentication** tab options (excluding the **Serial** option; serial access is allowed).

- Local users

Configure the **Access Restriction** option for a given username. By default, the access restriction is **Full Access**, which allows full access to any services specified by the **Authentication** tab options.

Configure 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.

After you configure the password policy, when you change a password (either your own or another user's), the password policy applies to the new password. Any existing passwords are grandfathered in. The new policy applies to changing the password with the **User Accounts** pane as well as the **Change My Password** pane.

Before You Begin

- Configure both CLI/ASDM and enable authentication.
- Specify the local database.

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > Password Policy**.

Step 2 Configure any mix of the following options:

- **Minimum Password Length**—Enter the minimum length for passwords. Valid values range from 3 to 64 characters. The recommended minimum password length is 8 characters.
- **Lifetime**—Enter the interval in days after which passwords expire for remote users (SSH, Telnet, HTTP); users at the console port are never locked out due to password expiration. Valid values are between 0 and 65536 days. The default value is 0 days, a value indicating that passwords will never expire.

7 days before the password expires, a warning message appears. After the password expires, system access is denied to remote users. To gain access after expiration, do one of the following:

- Have another administrator change your password.
- Log in to the physical console port to change your password.
- **Minimum Number Of**—Specify the minimum of characters from the following types:
 - **Numeric Characters**—Enter the minimum number of numeric characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0.
 - **Lower Case Characters**—Enter the minimum number of lower case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.
 - **Upper Case Characters**—Enter the minimum number of upper case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.

- **Special Characters**—Enter the minimum number of special characters that passwords must have. Valid values range from 0 to 64 characters. Special characters include the following: !, @, #, \$, %, ^, &, *, '(', and ')'. The default value is 0.
- **Different Characters from Previous Password**—Enter the minimum number of characters that you must change between new and old passwords. Valid values are between 0 and 64 characters. The default value is 0. Character matching is position independent, meaning that new password characters are considered changed only if they do not appear anywhere in the current password.

Step 3 (Optional) Check the **Authentication Enable** check box to require users to change their password on the **Change My Password** pane instead of the **User Accounts** pane. The default setting is disabled: a user can use either method to change their password.

If you enable this feature and try to change your password on the **User Accounts** pane, the following error message is generated:

ERROR: Changing your own password is prohibited

Step 4 Click **Apply** to save the configuration settings.

Change 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. If password policy authentication is not enabled, then you can use this method, or you can change your user account directly.

To change your username password, perform the following steps:

Procedure

- Step 1** Choose **Configuration > Device Management > Users/AAA > Change Password**.
- Step 2** Enter your old password.
- Step 3** Enter your new password.
- Step 4** Confirm your new password.
- Step 5** Click **Make Change**.
- Step 6** Click the **Save** icon to save your changes to the running configuration.
-

Configure 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

Configure 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).

To configure local command authorization, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Users/AAA > AAA Access > Authorization**.

Step 2 Check the **Enable authorization for command access > Enable** check box.

Step 3 Choose **LOCAL** from the **Server Group** drop-down list.

Step 4 When you enable local command authorization, you have the option of manually assigning privilege levels to individual commands or groups of commands or enabling the predefined user account privileges.

- Click **Set ASDM Defined User Roles** to use predefined user account privileges.

The **ASDM Defined User Roles Setup** dialog box appears. Click **Yes** to use the predefined user account privileges: **Admin** (privilege level 15, with full access to all CLI commands; **Read Only** (privilege level 5, with read-only access); and **Monitor Only** (privilege level 3, with access to the **Monitoring** section only).

- Click **Configure Command Privileges** to manually configure command levels.

The **Command Privileges Setup** dialog box appears. You can view all commands by choosing **All Modes** from the **Command Mode** drop-down list, or you can choose a configuration mode to view the commands available in that mode. For example, if you choose context, you can view all commands available in context configuration mode. If a command can be entered in user EXEC or privileged EXEC mode as well as configuration mode, and the command performs different actions in each mode, you can set the privilege level for these modes separately.

The **Variant** column displays show, clear, or cmd. You can set the privilege only for the show, clear, or configure form of the command. The configure form of the command is typically the form that causes a configuration change, either as the unmodified command (without the **show** or **clear** prefix) or as the **no** form.

To change the level of a command, double-click it or click **Edit**. You can set the level between 0 and 15. You can only configure the privilege level of the *main* command. For example, you can configure the level of all **aaa** commands, but not the level of the **aaa authentication** command and the **aaa authorization** command separately.

To change the level of all commands that appear, click **Select All**, then **Edit**.

Click **OK** to accept your changes.

Step 5 Check the **Perform authorization for exec shell access > Enable** check box to support administrative user privilege levels from RADIUS.

This option also enables management authorization for local, RADIUS, mapped LDAP, and TACACS+ users.

Without this option, the ASA only supports privilege levels for local database users and defaults all other types of users to level 15.

Step 6 Click **Apply**.

The authorization settings are assigned, and the changes are saved to the running configuration.

Configure Commands on the TACACS+ Server

You can configure commands on a Cisco Secure Access Control Server (ACS) TACACS+ server as a shared profile component, for a group, or for individual users. For third-party TACACS+ servers, see your server documentation for more information about command authorization support.

See the following guidelines for configuring commands in Cisco Secure ACS Version 3.1; many of these guidelines also apply to third-party servers:

- The ASA sends the commands to be authorized as shell commands, so configure the commands on the TACACS+ server as shell commands.



Note Cisco Secure ACS might include a command type called “pix-shell.” Do not use this type for ASA command authorization.

- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

For example, to allow the **show running-configuration aaa-server** command, add **show running-configuration** to the command field, and type **permit aaa-server** in the arguments field.

- You can permit all arguments of a command that you do not explicitly deny by checking the **Permit Unmatched Args** check box.

For example, you can configure just the **show** command, then all the **show** commands are allowed. We recommend using this method so that you do not have to anticipate every variant of a command, including abbreviations and a question mark, which shows CLI usage (see the following figure).

Figure 34-1 *Permitting All Related Commands*

- For commands that are a single word, you *must* permit unmatched arguments, even if there are no arguments for the command, for example **enable** or **help** (see the following figure).

Figure 34-2 *Permitting Single Word Commands*


- To disallow some arguments, enter the arguments preceded by **deny**.

For example, to allow **enable**, but not **enable password**, enter **enable** in the commands field, and **deny password** in the arguments field. Be sure to check the **Permit Unmatched Args** check box so that **enable** alone is still allowed (see the following figure).

Figure 34-3 *Disallowing Arguments*

- When you abbreviate a command at the command line, the ASA expands the prefix and main command to the full text, but it sends additional arguments to the TACACS+ server as you enter them.

For example, if you enter **sh log**, then the ASA sends the entire command to the TACACS+ server, **show logging**. However, if you enter **sh log mess**, then the ASA sends **show logging mess** to the TACACS+ server, and not the expanded command **show logging message**. You can configure multiple spellings of the same argument to anticipate abbreviations (see the following figure).

Figure 34-4 Specifying Abbreviations

- We recommend that you allow the following basic commands for all users:
 - **show checksum**
 - **show curpriv**
 - **enable**
 - **help**
 - **show history**
 - **login**
 - **logout**
 - **pager**
 - **show pager**
 - **clear pager**
 - **quit**
 - **show version**

Configure 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 [Recover from a Lockout, page 34-24](#).

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 [Configure Command Authorization](#),

[page 34-17](#).

To configure command authorization using a TACACS+ server, perform the following steps:

Procedure

-
Step 1	Choose Configuration > Device Management > Users/AAA > AAA Access > Authorization .
Step 2	Check the Enable authorization for command access > Enable check box.
Step 3	Choose a AAA server group name from the Server Group drop-down list.
Step 4	(Optional) You can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. To do so, check the Use LOCAL when server group fails check box. We recommend that you use the same username and password in the local database as the AAA server, because the ASA prompt does not give any indication which method is being used. Be sure to configure users in the local database and command privilege levels.
Step 5	Click Apply .
- The command authorization settings are assigned, and the changes are saved to the running configuration.
-

Configure 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:

Procedure

-
Step 1	To enable accounting of users when they enter the enable command, perform the following steps: a. Choose Configuration > Device Management > Users/AAA > AAA Access > Accounting, then check the Require accounting to allow accounting of user activity > Enable check box. b. Choose a RADIUS or TACACS+ server group name.
Step 2	To enable accounting of users when they access the ASA using Telnet, SSH, or the serial console, perform the following steps: a. Check the Serial, SSH, and/or Telnet check boxes in the Require accounting for the following types of connections area. b. Choose a RADIUS or TACACS+ server group name for each connection type.
Step 3	To configure command accounting, perform the following steps: a. Check the Enable check box in the Require accounting for the following types of connections area. b. Choose a TACACS+ server group name. RADIUS is not supported. You can send accounting messages to the TACACS+ accounting server when you enter any command other than show commands at the CLI.

- c. If you customize the command privilege level using the **Command Privilege Setup** dialog box, you can limit which commands the ASA accounts for by specifying a minimum privilege level in the **Privilege level** drop-down list. The ASA does not account for commands that are below the minimum privilege level.

Step 4 Click **Apply**.

The accounting settings are assigned, and the changes are saved to the running configuration.

Set a Management Session Quota

You can establish a maximum number of simultaneous ASDM, SSH, and Telnet sessions that are allowed on the ASA. If the maximum is reached, no additional sessions are allowed and a syslog message is generated. To prevent a system lockout, the management session quota mechanism cannot block a console session.

To set a management session quota, perform the following steps:

Procedure

Step 1 Choose **Configuration > Device Management > Management Access > Management Session Quota**.

Step 2 Enter the maximum number of simultaneous ASDM, SSH, and Telnet sessions that are allowed on the ASA. Valid values range from 0 to 10000.



Note If the management quota session number is exceeded, an error message appears, and ASDM closes.

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

Recover 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.

The following table lists the common lockout conditions and how you might recover from them.

Table 34-1 CLI Authentication and Command Authorization Lockout Scenarios

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
Local CLI authentication	No users have been configured in the local database.	If you have no users in the local database, you cannot log in, and you cannot add any users.	Log in and reset the passwords and aaa commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and add a user.
TACACS+ command authorization TACACS+ CLI authentication RADIUS CLI authentication	The server is down or unreachable and you do not have the fallback method configured.	If the server is unreachable, then you cannot log in or enter any commands.	1. Log in and reset the passwords and AAA commands. 2. Configure the local database as a fallback method so you do not get locked out when the server is down.	1. If the server is unreachable because the network configuration is incorrect on the ASA, session into the ASA from the switch. From the system execution space, you can change to the context and reconfigure your network settings. 2. Configure the local database as a fallback method so that you do not get locked out when the server is down.
TACACS+ command authorization	You are logged in as a user without enough privileges or as a user that does not exist.	You enable command authorization, but then find that the user cannot enter any more commands.	Fix the TACACS+ server user account. If you do not have access to the TACACS+ server and you need to configure the ASA immediately, then log into the maintenance partition and reset the passwords and aaa commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and complete the configuration changes. You can also disable command authorization until you fix the TACACS+ configuration.
Local command authorization	You are logged in as a user without enough privileges.	You enable command authorization, but then find that the user cannot enter any more commands.	Log in and reset the passwords and aaa commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and change the user level.

Monitoring Device Access

See the following screens for monitoring device access:

- **Monitoring > Properties > Device Access > ASDM/HTTPS/Telnet/SSH Sessions**

The top pane lists the connection types, session IDs, and IP addresses for users connected through ASDM, HTTPS, and Telnet sessions. To disconnect a specific session, click **Disconnect**.

The bottom pane lists the clients, usernames, connection states, software versions, incoming encryption types, outgoing encryption types, incoming HMACs, outgoing HMACs, SSH session IDs, remaining rekey data, remaining rekey time, data-based rekeys, time-based rekeys, and the last rekey time. To disconnect a specific session, click **Disconnect**.

- **Monitoring > Properties > Device Access > Authenticated Users**

This pane lists the usernames, IP addresses, dynamic ACLs, inactivity timeouts (if any), and absolute timeouts for users who were authenticated by AAA servers.

- **Monitoring > Properties > Device Access > AAA Local Locked Out Users**

This pane lists the usernames of locked-out AAA local users, the number of failed attempts to authenticate, and the times that users were locked out. To clear a specific user who has been locked out, click **Clear Selected Lockout**. To clear all users who have been locked out, click **Clear All Lockouts**.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

History for Management Access

Table 34-2 History for Management Access

Feature Name	Platform Releases	Description
Management Access	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH Configuration > Device Management > Management Access > Command Line (CLI) > Banner Configuration > Device Management > Management Access > CLI Prompt Configuration > Device Management > Management Access > ICMP Configuration > Device Management > Management Access > File Access > FTP Client Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server Configuration > Device Management > Management Access > File Access > Mount-Points Configuration > Device Management > Users/AAA > AAA Access > Authentication Configuration > Device Management > Users/AAA > AAA Access > Authorization Configuration > Device Management > Users/AAA > AAA Access > > Accounting.
Increased SSH security; the SSH default username is no longer supported.	8.4(2)	Starting in 8.4(2), you can no longer connect to the ASA using SSH with the pix or asa username and the login password. To use SSH, you must configure AAA authentication using the aaa authentication ssh console LOCAL command (CLI) or Configuration > Device Management > Users/AAA > AAA Access > Authentication (ASDM); then define a local user by entering the username command (CLI) or choosing Configuration > Device Management > Users/AAA > User Accounts (ASDM). If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.
Support for administrator password policy when using the local database	8.4(4.1), 9.1(2)	When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters. We introduced the following screen: Configuration > Device Management > Users/AAA > Password Policy.

Table 34-2 History for Management Access (continued)

Feature Name	Platform Releases	Description
Support for SSH public key authentication	8.4(4.1), 9.1(2)	You can enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF. <i>PKF key format support is only in 9.1(2) and later.</i>
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	8.4(4.1), 9.1(2)	Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported. We modified the following screen: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH.
Support for a maximum number of management sessions	8.4(4.1), 9.1(2)	You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions. We introduced the following screen: Configuration > Device Management > Management Access > Management Session Quota.
For the ASASM in multiple context mode, support for Telnet and virtual console authentication from the switch.	8.5(1)	Although connecting to the ASASM from the switch in multiple context mode connects to the system execution space, you can configure authentication in the admin context to govern those connections.
AES-CTR encryption for SSH	9.1(2)	The SSH server implementation in the ASA now supports AES-CTR mode encryption.
Improved SSH rekey interval		An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic. .
Improved one-time password authentication	9.2(1)	Administrators who have sufficient authorization privileges may enter privileged EXEC mode by entering their authentication credentials once. The auto-enable option was added to the aaa authorization exec command. We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization.

Table 34-2 History for Management Access (continued)

Feature Name	Platform Releases	Description
ASDM Username From Certificate Configuration (DoD)	9.4(1)	This feature introduces the ability to authorize users by extracting the username from the certificate along with using the username provided by the user. We added the following screen: Configuration > Device Management > Management Access > HTTP Certificate Rule. We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization.



Software and Configurations

This chapter describes how to manage the Cisco ASA software and configurations.

- [Upgrade the Software, page 35-1](#)
- [Manage Files, page 35-8](#)
- [Set the ASA Image, ASDM, and Startup Configuration, page 35-15](#)
- [Back Up and Restore Configurations or Other Files, page 35-19](#)
- [Save the Running Configuration to a TFTP Server, page 35-23](#)
- [Schedule a System Restart, page 35-24](#)
- [Downgrade Your Software, page 35-24](#)
- [Configure Auto Update, page 35-26](#)
- [History for Software and Configurations, page 35-31](#)

Upgrade the Software

This section describes how to upgrade your single device, failover devices, or clustering devices.

- [Upgrade Path, page 35-1](#)
- [View Your Current Version, page 35-2](#)
- [Download the Software from Cisco.com, page 35-2](#)
- [Upgrade a Standalone Unit, page 35-2](#)
- [Upgrade a Failover Pair or ASA Cluster, page 35-4](#)

Upgrade Path

See the following table for the upgrade path for your version. Some versions require an interim upgrade before you can upgrade to the latest version.

Note: There are no special requirements for Zero Downtime Upgrades for failover and ASA clustering with the following exception. Upgrading ASA clustering from 9.0(1) or 9.1(1): due to CSCue72961, hitless upgrading is not supported.

Current ASA Version	First Upgrade to:	Then Upgrade to:
8.2(x) and earlier	8.4(6)	9.4(1) or later
8.3(x)	8.4(6)	9.4(1) or later
8.4(1) through 8.4(4)	8.4(6), 9.0(4), or 9.1(2)	9.4(1) or later
8.4(5) and later	—	9.4(1) or later
8.5(1)	9.0(4) or 9.1(2)	9.4(1) or later
8.6(1)	9.0(4) or 9.1(2)	9.4(1) or later
9.0(1)	9.0(4) or 9.1(2)	9.4(1) or later
9.0(2) or later	—	9.4(1) or later
9.1(1)	9.1(2)	9.4(1) or later
9.1(2) or later	—	9.4(1) or later
9.2(x)	—	9.4(1) or later
9.3(x)	—	9.4(1) or later

Configuration Migration

Depending on your current version, you might experience one or more configuration migrations when you upgrade. For example, when upgrading from 8.0 to 9.4, you will experience all of these migrations:

- 8.2—See the [8.2 release notes](#).
- 8.3—See the [Cisco ASA 5500 Migration Guide to Version 8.3](#).
- 8.4—See the [8.4 upgrade guide](#).
- 9.0—See the [9.0 upgrade guide](#).

View Your Current Version

The software version appears on the ASDM home page; view the home page to verify the software version of your ASA.

Download the Software from Cisco.com

If you are using the ASDM Upgrade Wizard, you do not have to pre-download the software. If you are manually upgrading, for example for a failover upgrade, download the images to your local computer.

If you have a Cisco.com login, you can obtain the OS and ASDM images from the following website:

<http://www.cisco.com/go/asa-software>

Upgrade a Standalone Unit

This section describes how to install the ASDM and operating system (OS) images.

- [Upgrade from Your Local Computer, page 35-3](#)
- [Upgrade Using the Cisco.com Wizard, page 35-3](#)

Upgrade from Your Local Computer

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

Procedure

1. (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.
2. In the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**. The **Upgrade Software** dialog box appears.
3. From the **Image to Upload** drop-down list, choose **ASDM**.
4. In the **Local File Path** field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.
5. In the **Flash File System Path** field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.
6. Click **Upload Image**. The uploading process might take a few minutes.
7. You are prompted to set this image as the ASDM image. Click **Yes**.
8. You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the **Upgrade** tool.
Note: You will save the configuration and reload ASDM *after* you upgrade the ASA software.
9. Repeat 2. through 8., choosing **ASA** from the **Image to Upload** drop-down list. You can also use this procedure to upload other file types.
10. Choose **Tools > System Reload** to reload the ASA.
A new window appears that asks you to verify the details of the reload.
 - a. Click the **Save the running configuration at the time of reload** radio button (the default).
 - b. Choose a time to reload (for example, **Now**, the default).
 - c. Click **Schedule Reload**.
Once the reload is in progress, a **Reload Status** window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.
11. After the ASA reloads, restart ASDM.

Upgrade Using the Cisco.com Wizard

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

In this wizard, you can do the following:

- Choose an ASA image file and/or ASDM image file to upgrade.
Note: ASDM downloads the latest image version, which includes the build number. For example, if you are downloading 9.4(1), the download might be 9.4(1.2). This behavior is expected, so you may proceed with the planned upgrade.
- Review the upgrade changes that you have made.
- Download the image or images and install them.
- Review the status of the installation.

- If the installation completed successfully, restart the ASA to save the configuration and complete the upgrade.

Procedure

1. (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.
2. Choose **Tools > Check for ASA/ASDM Updates**.
In multiple context mode, access this menu from the System.
The **Cisco.com Authentication** dialog box appears.
3. Enter your Cisco.com username and password, and then click **Login**.
The **Cisco.com Upgrade Wizard** appears.
Note: If there is no upgrade available, a dialog box appears. Click **OK** to exit the wizard.
4. Click **Next** to display the **Select Software** screen.
The current ASA version and ASDM version appear.
5. To upgrade the ASA version and ASDM version, perform the following steps:
 - a. In the **ASA** area, check the **Upgrade to** check box, and then choose an ASA version to which you want to upgrade from the drop-down list.
 - b. In the **ASDM** area, check the **Upgrade to** check box, and then choose an ASDM version to which you want to upgrade from the drop-down list.
6. Click **Next** to display the **Review Changes** screen.
7. Verify the following items:
 - The ASA image file and/or ASDM image file that you have downloaded are the correct ones.
 - The ASA image file and/or ASDM image file that you want to upload are the correct ones.
 - The correct ASA boot image has been selected.
8. Click **Next** to start the upgrade installation.
You can then view the status of the upgrade installation as it progresses.
The **Results** screen appears, which provides additional details, such as the upgrade installation status (success or failure).
9. If the upgrade installation succeeded, for the upgrade versions to take effect, check the **Save configuration and reload device now** check box to restart the ASA, and restart ASDM.
10. Click **Finish** to exit the wizard and save the configuration changes that you have made.
Note: To upgrade to the next higher version, if any, you must restart the wizard.

Upgrade a Failover Pair or ASA Cluster

To perform a zero downtime upgrade, you need to upgrade each unit in a particular order.

- [Upgrade an Active/Standby Failover Pair, page 35-5](#)
- [Upgrade an Active/Active Failover Pair, page 35-6](#)
- [Upgrade an ASA Cluster, page 35-7](#)

Upgrade an Active/Standby Failover Pair

To upgrade the Active/Standby failover pair, perform the following steps.

Procedure

1. (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.
2. On the active unit, in the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.
The **Upgrade Software** dialog box appears.
3. From the **Image to Upload** drop-down list, choose **ASDM**.
4. In the **Local File Path** field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.
5. In the **Flash File System Path** field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.
6. Click **Upload Image**. The uploading process might take a few minutes.
7. You are prompted to set this image as the ASDM image. Click **Yes**.
8. You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the **Upgrade** tool.
Note: You will save the configuration and reload ASDM *after* you upgrade the ASA software.
9. Repeat 2. through 8., choosing **ASA** from the **Image to Upload** drop-down list.
10. Click the **Save** icon on the toolbar to save your configuration changes.
11. Connect ASDM to the *standby* unit, and upload the ASA and ASDM software according to 2. through 9., using the same file locations you used on the active unit.
12. Choose **Tools > System Reload** to reload the standby ASA.

A new window appears that asks you to verify the details of the reload.

- a. Click the **Save the running configuration at the time of reload** radio button (the default).
- b. Choose a time to reload (for example, **Now**, the default).
- c. Click **Schedule Reload**.

Once the reload is in progress, a **Reload Status** window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

13. After the standby ASA reloads, restart ASDM and connect to the standby unit to make sure it is running.
14. Connect ASDM to the *active* unit again.
15. Force the active unit to fail over to the standby unit by choosing **Monitoring > Properties > Failover > Status**, and clicking **Make Standby**.
16. Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload.

- a. Click the **Save the running configuration at the time of reload** radio button (the default).
- b. Choose a time to reload (for example, **Now**, the default).
- c. Click **Schedule Reload**.

Once the reload is in progress, a **Reload Status** window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

After the ASA comes up, it will now be the standby unit.

Upgrade an Active/Active Failover Pair

To upgrade two units in an Active/Active failover configuration, perform the following steps.

Before You Begin

Perform these steps in the system execution space. .

Procedure

1. (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.
2. On the primary unit, in the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.
The **Upgrade Software** dialog box appears.
3. From the **Image to Upload** drop-down list, choose **ASDM**.
4. In the **Local File Path** field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.
5. In the **Flash File System Path** field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.
6. Click **Upload Image**. The uploading process might take a few minutes.
7. You are prompted to set this image as the ASDM image. Click **Yes**.
8. You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the **Upgrade** tool.
Note: You will save the configuration and reload ASDM *after* you upgrade the ASA software.
9. Repeat 2. through 8., choosing **ASA** from the **Image to Upload** drop-down list.
10. Click the **Save** icon on the toolbar to save your configuration changes.
11. Make both failover groups active on the primary unit by choosing **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to move to the primary unit, and clicking **Make Active**.
12. Connect ASDM to the *secondary* unit, and upload the ASA and ASDM software according to 2. through 9., using the same file locations you used on the active unit.
13. Choose **Tools > System Reload** to reload the secondary ASA.
A new window appears that asks you to verify the details of the reload.
 - a. Click the **Save the running configuration at the time of reload** radio button (the default).
 - b. Choose a time to reload (for example, **Now**, the default).
 - c. Click **Schedule Reload**.Once the reload is in progress, a **Reload Status** window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.
14. Connect ASDM to the *primary* unit, and check when the secondary unit reloads by choosing **Monitoring > Failover > System**.
15. After the secondary unit comes up, force the primary unit to fail over to the secondary unit by choosing **Monitoring > Properties > Failover > System**, and clicking **Make Standby**.
16. Choose **Tools > System Reload** to reload the (formerly) active ASA.

A new window appears that asks you to verify the details of the reload.

- a. Click the **Save the running configuration at the time of reload** radio button (the default).
- b. Choose a time to reload (for example, **Now**, the default).
- c. Click **Schedule Reload**.

Once the reload is in progress, a **Reload Status** window appears that indicates that a reload is being performed. An option to exit ASDM is also provided.

If the failover groups are configured with Preempt Enabled, they automatically become active on their designated unit after the preempt delay has passed. If the failover groups are not configured with Preempt Enabled, you can return them to active status on their designated units using the **Monitoring > Failover > Failover Group #** pane.

Upgrade 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.

Procedure

1. Launch ASDM on the master unit.
2. (If there is a configuration migration) In ASDM, back up your existing configuration using the **Tools > Backup Configurations** tool.
3. In the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.
The **Upgrade Software from Local Computer** dialog box appears.
4. Click the **All devices in the cluster** radio button.
The **Upgrade Software** dialog box appears.
5. From the **Image to Upload** drop-down list, choose **ASDM**.
6. In the **Local File Path** field, enter the local path to the file on your computer or click **Browse Local Files** to find the file on your PC.
7. In the **Flash File System Path** field, enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.
8. Click **Upload Image**. The uploading process might take a few minutes.
9. You are prompted to set this image as the ASDM image. Click **Yes**.
10. You are reminded to exit ASDM and save the configuration. Click **OK**. You exit the Upgrade tool.
Note: You will save the configuration and reload ASDM *after* you upgrade the ASA software.
11. Repeat 3. through 10., choosing **ASA** from the **Image to Upload** drop-down list.
12. Click the **Save** icon on the toolbar to save your configuration changes.
13. Choose **Tools > System Reload**.
The System Reload dialog box appears.
14. Reload each slave unit one at a time by choosing a slave unit name from the Device drop-down list, and then clicking **Schedule Reload** to reload the unit now.

To avoid connection loss and allow traffic to stabilize, wait for each unit to come back up (approximately 5 minutes) before reloading the next unit. To view when a unit rejoins the cluster, see the **Monitoring > ASA Cluster > Cluster Summary** pane.

15. After all slave units have reloaded, disable clustering on the master unit by choosing **Configuration > Device Management > High Availability and Scalability > ASA Cluster**, uncheck the **Participate in ASA cluster** check box, and click **Apply**.

Wait for 5 minutes for a new master to be selected and traffic to stabilize. When the former master unit rejoins the cluster, it will be a slave.

Do not save the configuration; when the master unit reloads, you want clustering to be enabled on it.

16. Choose **Tools > System Reload** and reload the master unit from the System Reload dialog box by choosing **--This Device--** from the Device drop-down list.
 17. Quit and restart ASDM; you will reconnect to the new master unit.
-

Manage Files

ASDM provides a set of file management tools to help you perform basic file management tasks. The File Management tool lets you view, move, copy, and delete files stored in flash memory, transfer files, and to manage files on remote storage devices (mount points).

**Note**

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

- [Configure File Access, page 35-8](#)
- [Access the File Management Tool, page 35-12](#)
- [Transfer Files, page 35-13](#)

Configure File Access

The ASA can use an FTP client, secure copy client, or TFTP client. You can also configure the ASA as a secure copy server so you can use a secure copy client on your computer.

- [Configure the FTP Client Mode, page 35-9](#)
- [Configure the ASA as a Secure Copy Server, page 35-9](#)
- [Configure the ASA TFTP Client Path, page 35-10](#)
- [Add Mount Points, page 35-11](#)

Configure 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.

Procedure

-
- Step 1** From the Configuration > Device Management > Management Access > File Access > FTP Client pane, check the **Specify FTP mode as passive** check box.
- Step 2** Click **Apply**.
- The FTP client configuration is changed and the change is saved to the running configuration.
-

Configure 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.

Before You Begin

- 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 or wildcards.
- Enable SSH on the ASA according to [Configure ASA Access for ASDM, Telnet, or SSH, page 34-4](#).
- The ASA license must have the strong encryption (3DES/AES) license to support SSH Version 2 connections.
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

-
- Step 1** Depending on your context mode:
- For single mode, choose **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP)**.
 - For multiple mode in the System, choose **Configuration > Device Management > Device Administration > Secure Copy**
- Step 2** Check the **Enable secure copy server** check box.
- Step 3** (Optional) The ASA stores the SSH host key for each SCP server to which it connects. You can manually add or delete servers and their keys from the ASA database if desired.
- To add a key:
- a. Click **Add** for a new server, or select the server from the Trusted SSH Hosts table, and click **Edit**.
 - b. For a new server, in the Host field, enter the server IP address.

- c. Check the **Add public key for the trusted SSH host** check box.
- d. Specify one of the following keys:
 - Fingerprint—Enter the already hashed key; for example, a key that you copied from **show** command output.
 - Key—Enter the public key or hashed value of the SSH host. The key string is the Base64 encoded RSA public key of the remote peer. You can obtain the public key value from an open SSH client; that is, from the `.ssh/id_rsa.pub` file. After you submit the Base64 encoded public key, that key is then hashed via SHA-256.

To delete a key:

- a. Select the server from the Trusted SSH Hosts table, and click **Delete**.

Step 4 (Optional) To be informed when a new host key is detected, check the **Inform me when a new host key is detected** check box.

By default, this option is enabled. When this option is enabled, you are prompted to accept or reject the host key if it is not already stored on the ASA. When this option is disabled, the ASA accepts the host key automatically if it was not stored before.

Step 5 Click **Apply**.

Examples

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.

Configure 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*. In this way, you can back up and propagate configuration files to multiple ASAs.

This section lets you predefine the path to a TFTP server so you do not need to enter it in commands such as **copy** and **configure net**.

Procedure

- Step 1** Choose **Configuration > Device Management > Management Access > File Access > TFTP Client**, and check the **Enable** check box.
- Step 2** From the Interface Name drop-down list, choose the interface to use as a TFTP client.
- Step 3** In the IP Address field, enter the IP address of the TFTP server on which configuration files will be saved.
- Step 4** In the Path field, enter the path to the TFTP server on which configuration files will be saved.
For example: `/tftpboot/asa/config3`
- Step 5** Click **Apply**.

Related Topics

[Transfer Files, page 35-13](#)

Add Mount Points

You can add a CIFS or FTP mount point.

- [Add a CIFS Mount Point, page 35-11](#)
- [Add an FTP Mount Point, page 35-11](#)

Add a CIFS Mount Point

To define a Common Internet File System (CIFS) mount point, perform the following steps.

Procedure

-
Step 1	Choose Configuration > Device Management > Management Access > File Access > Mount-Points , and click Add > CIFS Mount Point . The Add CIFS Mount Point dialog box appears.
Step 2	Check the Enable mount point check box. This option attaches the CIFS file system on the ASA to the UNIX file tree.
Step 3	In the Mount Point Name field, enter the name of an existing CIFS location.
Step 4	In the Server Name or IP Address field, enter the name or IP address of the server in which the mount point is located.
Step 5	In the Share Name field, enter the name of the folder on the CIFS server.
Step 6	In the NT Domain Name field, enter the name of the NT Domain in which the server resides.
Step 7	In the User Name field, enter the name of the user authorized for file system mounting on the server.
Step 8	In the Password field, enter the password for the user authorized for file system mounting on the server.
Step 9	In the Confirm Password field, reenter the password.
Step 10	Click OK . The Add CIFS Mount Point dialog box closes.
Step 11	Click Apply .
-

Add an FTP Mount Point

For an FTP mount point, the FTP server must have a UNIX directory listing style. Microsoft FTP servers have a default of the MS-DOS directory listing style.

Procedure

-
Step 1	Choose Configuration > Device Management > Management Access > File Access > Mount-Points , and click Add > FTP Mount Point . The Add FTP Mount Point dialog box appears.

- Step 2** Check the **Enable** check box.
This option attaches the FTP file system on the ASA to the UNIX file tree.
- Step 3** In the Mount Point Name field, enter the name of an existing FTP location.
- Step 4** In the Server Name or IP Address field, enter the name or IP address of the server where the mount point is located.
- Step 5** In the Mode field, click the radio button for the FTP mode (**Active** or **Passive**). When you choose Passive mode, the client initiates both the FTP control connection and the data connection. The server responds with the number of its listening port for this connection.
- Step 6** In the Path to Mount field, enter the directory path name to the FTP file server.
- Step 7** In the User Name field, enter the name of the user authorized for file system mounting on the server.
- Step 8** In the Password field, enter the password for the user authorized for file system mounting on the server.
- Step 9** In the Confirm Password field, reenter the password.
- Step 10** Click **OK**.
The Add FTP Mount Point dialog box closes.
- Step 11** Click **Apply**.
-

Access the File Management Tool

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

Procedure

-
- Step 1** In the main ASDM application window, choose **Tools > File Management**.
The File Management dialog box appears.
- The Folders pane displays the available folders on disk.
 - Flash Space shows the total amount of flash memory and how much memory is available.
 - The Files area displays the following information about files in the selected folder:
 - Path
 - Filename
 - Size (bytes)
 - Time Modified
 - Status, which indicates whether a selected file is designated as a boot configuration file, boot image file, ASDM image file, SVC image file, CSD image file, or APCF image file.
- Step 2** Click **View** to display the selected file in your browser.
- Step 3** Click **Cut** to cut the selected file for pasting to another directory.
- Step 4** Click **Copy** to copy the selected file for pasting to another directory.
- Step 5** Click **Paste** to paste the copied file to the selected destination.
- Step 6** Click **Delete** to remove the selected file from flash memory.
- Step 7** Click **Rename** to rename a file.

- Step 8** Click **New Directory** to create a new directory for storing files.
- Step 9** Click **File Transfer** to open the File Transfer dialog box. See [Transfer Files, page 35-13](#) for more information.
- Step 10** Click **Mount Points** to open the Manage Mount Points dialog box. See [Add Mount Points, page 35-11](#) for more information.
-

Transfer Files

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



Note

For the IPS SSP software module, before you download the IPS software to disk0, make sure at least 50% of the flash memory is free. When you install IPS, IPS reserves 50% of the internal flash memory for its file system.

- [Transfer Files Between Local PC and Flash, page 35-13](#)
- [Transfer Files Between Remote Server and Flash, page 35-13](#)

Transfer Files Between Local PC and Flash

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

Procedure

- Step 1** In the main ASDM application window, choose **Tools > File Management**.
The File Management dialog box appears.
- Step 2** Click the down arrow next to **File Transfer**, and then click **Between Local PC and Flash**.
The File Transfer dialog box appears.
- Step 3** Select and *drag* the file(s) from either your local computer or the flash file system that you want to upload or download to the desired location. Alternatively, select the file(s) from either your local computer or the flash file system that you want to upload or download, and click the right arrow or left arrow to transfer the file(s) to the desired location.
- Step 4** Click **Close** when you are done.
-

Transfer Files Between Remote Server and Flash

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

Procedure

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

The File Management dialog box appears.

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

The File Transfer dialog box appears.

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

- Step 4** Define the source file to be transferred.

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



Note File transfer supports IPv4 and IPv6 addresses.

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

- ap—ASCII files in passive mode
- an—ASCII files in non-passive mode
- ip—Binary image files in passive mode
- in—Binary image files in non-passive mode

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

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

- Step 7** In addition, you can copy a file from your startup configuration, running configuration, or an SMB file system through the CLI. For instructions about using the **copy** command, see the CLI configuration guide.

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

- a. To transfer the file to the flash file system, choose the **Flash file system** option.
- b. Enter the path to the location of the file or click **Browse Flash** to find the file location.

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

- a. Enter the path to the location of the file.
- b. For FTP transfers, enter the type. Valid types are the following:
- ap—ASCII files in passive mode
 - an—ASCII files in non-passive mode
 - ip—Binary image files in passive mode
 - in—Binary image files in non-passive mode

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

The Enter Username and Password dialog box appears.

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

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

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

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

Set the ASA Image, ASDM, and Startup Configuration

If you have more than one ASA or ASDM image, you should specify the image that you want to boot. If you do not set the image, the default boot image is used, and that image may not be the one intended. For the startup configuration, you can optionally specify a configuration file.

See the following default settings:

- **ASA Image:**
 - **Physical ASA**—Boots the first application image that it finds in internal flash memory.
 - **ASAv**—Boots the image in the read-only boot:/ partition that was created when you first deployed. You can upgrade the image in flash memory and configure the ASAv to boot from that image. Note that if you later clear your configuration, then the ASAv will revert to loading the original deployment image.
- **ASDM Image on All ASAs**—Boots the first ASDM image that it finds in internal flash memory, or if one does not exist in this location, then in external flash memory.
- **Startup Configuration**—By default, the ASA boots from a startup configuration that is a hidden file.

Procedure

-
Step 1	Choose Configuration > Device Management > System Image/Configuration > Boot Image/Configuration.
	You can specify up to four local binary image files for use as the startup image, and one image located on a TFTP server for the device to boot from. If you specify an image located on a TFTP server, it must be first in the list. If the device cannot reach the TFTP server to load the image, it tries to load the next image file in the list located in flash.
Step 2	Click Add in the Boot Image/Configuration pane.
Step 3	Browse to the image from which you want to boot. For a TFTP image, enter the TFTP URL in the File Name field. Click OK .
Step 4	Arrange the images in order by using the Move Up and Move Down buttons.
Step 5	(Optional) In the Boot Configuration File Path field, specify the startup configuration file by clicking Browse Flash and choosing the configuration. Click OK .
Step 6	In the ASDM Image File Path field, specify the ASDM image by clicking Browse Flash and choosing the image. Click OK .
Step 7	Click Apply .
-

Use the ROM Monitor to Load an Image

You can load a new image using the ROM Monitor.

- [Load an Image for the ASA 5500-X Series Using ROM Monitor, page 35-16](#)
- [Load an Image for the ASASM Using ROM Monitor, page 35-17](#)
- [Recover and Load an Image for the ASA 5506W-X Wireless Access Point, page 35-18](#)

Load an Image for the ASA 5500-X Series Using ROM Monitor

To load a software image onto 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 [Access the Appliance Console, page 2-1](#).
- 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=tftp-home/asa941-smp-k9.bin
rommon #5> PORT=GigabitEthernet0/0
GigabitEthernet0/0
Link is UP
MAC Address: 0012.d949.15b8
```



Note Be sure that the connection to the network already exists.



Note The **PORT** command is ignored on the ASA 5506-X, ASA 5508-X, and ASA 5516-X models, and you must perform TFTP recovery on these platforms from the Management 1/1 interface.

- Step 5** Validate your settings:

```
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:

```
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:

```
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** Verify that the correct software image has been loaded into the ASA by checking the version in the ASA:
- ```
ciscoasa# show version
```

## Load an Image for the ASASM Using ROM Monitor

To load a software image to an ASASM from the ROM monitor mode using TFTP, perform the following steps.

### Procedure

- Step 1** Connect to the ASA console port according to the instructions in [Access the ASA Services Module Console, page 2-2](#).
- 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=tftp-main/asa941-smp-k9.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** Validate your settings:

```

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:

```

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:

```

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** Verify that the correct software image has been loaded into the ASA by checking the version in the ASA:

```

ciscoasa# show version

```

## Recover and Load an Image for the ASA 5506W-X Wireless Access Point

To recover and load a software image onto an ASA 5506W-X using TFTP, perform the following steps.

**Procedure****Step 1** Session to the access point (AP) and enter the AP ROMMON (not the ASA ROMMON):

```

ciscoasa# hw-module module wlan recover image

```

**Step 2** Follow the procedure in the [Cisco IOS Software Configuration Guide for Cisco Aironet Access Points](#).

---

## Back Up and Restore Configurations or Other Files

We recommend that you make regular backups of your configuration and other system files to guard against system failure.

- [Perform a Complete System Backup or Restoration, page 35-19](#)
- 
- [Back Up the Local CA Server, page 35-22](#)
- [Save the Running Configuration to a TFTP Server, page 35-23](#)

## Perform a Complete System Backup or Restoration

These procedures describe how to back up and restore configurations and images to a zip file and transfer it to your local computer.

- [Before You Begin, page 35-19](#)
- [Back Up the System, page 35-20](#)
- [Restore the Backup, page 35-22](#)

### Before You Begin

- You should have at least 300 MB of disk space available at the backup or restore location before you start a backup or restore.
- The ASA must be in single context mode.
- If you make any configuration changes during or after a backup, those changes will not be included in the backup. If you change a configuration after making the backup, then perform a restore, this configuration change will be overwritten. As a result, the ASA might behave differently.
- You can start only one backup or restore at a time.
- You can only restore a configuration to the same ASA version as when you performed the original backup. You cannot use the restore tool to migrate a configuration from one ASA version to another. If a configuration migration is required, the ASA automatically upgrades the resident startup configuration when it loads the new ASA OS.
- If you use clustering, you can only back up or restore the startup-configuration, running-configuration, and identity certificates. You must create and restore a backup separately for each unit.
- If you use failover, you must create and restore a backup separately for the active and standby units.
- If you set a master passphrase for the ASA, then you need that master passphrase to restore the backup configuration that you create with this procedure. If you do not know the master passphrase for the ASA, see [Configure the Master Passphrase, page 18-8](#) to learn how to reset it before continuing with the backup.

- If you import PKCS12 data (with the **crypto ca trustpoint** command) and the trustpoint uses RSA keys, the imported key pair is assigned the same name as the trustpoint. Because of this limitation, if you specify a different name for the trustpoint and its key pair after you have restored an ASDM configuration, the startup configuration will be the same as the original configuration, but the running configuration will include a different key pair name. This means that if you use different names for the key pair and trustpoint, you cannot restore the original configuration. To work around this issue, make sure that you use the same name for the trustpoint and its key pair.
- You cannot back up using the CLI and restore using ASDM, or vice versa.
- Each backup file includes the following content:
  - Running-configuration
  - Startup-configuration
  - All security images
    - Cisco Secure Desktop and Host Scan images
    - Cisco Secure Desktop and Host Scan settings
    - AnyConnect (SVC) client images and profiles
    - AnyConnect (SVC) customizations and transforms
  - Identity certificates (includes RSA key pairs tied to identity certificates; excludes standalone keys)
  - VPN pre-shared keys
  - SSL VPN configurations
  - Application Profile Custom Framework (APCF)
  - Bookmarks
  - Customizations
  - Dynamic Access Policy (DAP)
  - Plug-ins
  - Pre-fill scripts for connection profiles
  - Proxy Auto-config
  - Translation table
  - Web content
  - Version information

## Back Up the System

This procedure describes how to perform a complete system backup.

### Procedure

- 
- Step 1** Create a folder on your computer to store backup files so they will be easy to find in case you need to restore them later.
- Step 2** Choose **Tools > Backup Configurations**.



The Backup Configurations dialog box appears. Click the down arrow in the **SSL VPN Configuration** area to view the backup options for SSL VPN configurations. By default, all configuration files are checked and will be backed up if they are available. If you want to back up all of the files in the list, go to Step 5.

- Step 3** Uncheck the **Backup All** check box if you want to select the configurations to back up.
- Step 4** Check the check box next to the option that you want to back up.
- Step 5** Click **Browse Local to specify a directory and file name for the backup .zip file**.
- Step 6** In the Select dialog box, choose the directory in which you want to store the backup file.
- Step 7** Click **Select**. The path appears in the Backup File field.
- Step 8** Enter the name of the destination backup file after the directory path. The backup file name must be between 3 and 232 characters long.
- Step 9** Click **Backup**. The backup proceeds immediately unless you are backing up certificates or the ASA is using a master passphrase.
- Step 10** If you have configured and enabled a master passphrase on your ASA, you receive a warning message with a suggestion to change the master passphrase, if you do not know it, before proceeding with the backup. Click **Yes** to proceed with the backup if you know the master passphrase. The backup proceeds immediately unless you are backing up identity certificates.
- Step 11** If you are backing up an identity certificate, you are asked to enter a separate passphrase to be used for encoding the certificates in PKCS12 format. You can enter a passphrase or skip this step.



**Note** Identify certificates are backed up by this process; however, certificate authority certificates are not backed up. For instructions on backing up CA certificates, see [Back Up the Local CA Server, page 35-22](#).

- To encrypt certificates, enter and confirm your certificate passphrase in the Certificate Passphrase dialog box and click **OK**. You will need to remember the password you enter in this dialog box when restoring the certificates.
- Clicking **Cancel** skips the step and does not back up certificates.

After clicking OK or Cancel, the backup begins immediately.

- Step 12** After the backup is complete, the status window closes and the Backup Statistics dialog box appears to provide success and failure messages.



**Note** Backup “failure messages” are most likely caused by the lack of an existing configuration for the types indicated.

- Step 13** Click **OK** to close the Backup Statistics dialog box.

## Restore the Backup

You can specify configurations and images to restore from a zip tar.gz file on your local computer.

### Procedure

- 
- Step 1** Choose **Tools > Restore Configurations**.
- Step 2** In the Restore Configurations dialog box, click **Browse Local Directory**, choose the zip file on your local computer that contains the configuration to restore, then click **Select**. The path and the zip filename appear in the **Local File** field.
- The zip file that you restore must be created by choosing the **Tools > Backup Configurations** option.
- Step 3** Click **Next**. The second Restore Configuration dialog box appears. Check the check boxes next to the configurations that you want to restore. All available SSL VPN configurations are selected by default.
- Step 4** Click **Restore**.
- Step 5** If you specified a certificate passphrase with which to encrypt the certificates when you created the backup file, ASDM prompts you to enter the passphrase.
- Step 6** If you chose to restore the running configuration, you are asked if you want to merge the running configuration, replace the running configuration, or skip this part of the restoration process.
- Merging configurations combines the current running configuration and the backed-up running configuration.
  - Replacing the running configuration uses the backed-up running configuration only.
  - Skipping the step does not restore the backed-up running configuration.
- ASDM displays a status dialog box until the restore operation is finished.
- Step 7** If you replaced or merged the running configuration, close ASDM and restart it. If you did not restore the running configuration or the running configuration, refresh the ASDM session for the changes to take effect.
- 

## Back Up the Local CA Server

When you do an ASDM backup, it does not include the local CA server database, so you are not backing up the CA certificates stored on the server. If you want to back up the local CA server, use this manual process with the ASA CLI.

### Procedure

- 
- Step 1** Enter the **show run crypto ca server** command.

```
crypto ca server
keysize server 2048
subject-name-default OU=aa,O=Cisco,ST=ca,
issuer-name CN=xxx,OU=yyy,O=Cisco,L=Bxb,St=Mass
smtp from-address abcd@cisco.com
publish-crl inside 80
publish-crl outside 80
```

- Step 2** Use the **crypto ca import** command to import the local CA PKCS12 file to create the LOCAL-CA-SERVER trustpoint and to restore the keypair.

```
crypto ca import LOCAL-CA-SERVER pkcs12 <passphrase> (paste the pkcs12
base64 data here)
```



**Note** Be sure to use the exact name “LOCAL-CA-SERVER” for this step.

- Step 3** If the LOCAL-CA-SERVER directory does not exist, you need to create it by entering **mkdir LOCAL-CA-SERVER**.

- Step 4** Copy the local CA files into the LOCAL-CA-SERVER directory.

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.ser
disk0:/LOCAL-CA-SERVER/
```

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.cdb
disk0:/LOCAL-CA-SERVER/
```

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.udb
disk0:/LOCAL-CA-SERVER/
```

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.crl
disk0:/LOCAL-CA-SERVER/
```

```
copy ftp://10.10.1.1/CA-backup/LOCAL-CA-SERVER.p12
disk0:/LOCAL-CA-SERVER/
```

- Step 5** Enter the **crypto ca server** command to enable the local CA server

```
crypto ca server
no shutdown
```

- Step 6** Enter the **show crypto ca server** command to check that the local CA server is up and running.

- Step 7** Save the configuration.

## Save the Running Configuration to a TFTP Server

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

### Procedure

- Step 1** Choose **File > Save Running Configuration to TFTP Server**.

The Save Running Configuration to TFTP Server dialog box appears.

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

**Note**

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

## Schedule a System Restart

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

### Procedure

**Step 1** Choose **Tools > System Reload**.

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

- a. For the Configuration State, choose either to save or discard the running configuration at restart time.
- b. For the Reload Start Time, choose from the following options:
  - Click **Now** to perform an immediate restart.
  - Click **Delay by** to delay the restart by a specified amount of time. Enter the time before the restart begins in hours and minutes or only minutes.
  - Click **Schedule at** to schedule the restart to occur at a specific time and date. Enter the time of day the restart is to occur, and select the date of the scheduled restart.
- c. In the Reload Message field, enter a message to send to open instances of ASDM at restart time.
- d. Check the **On reload failure force immediate reload after** check box to show the amount of time elapsed in hours and minutes or only minutes before a restart is attempted again.
- e. Click **Schedule Reload** to schedule the restart as configured.

The Reload Status area displays the status of the restart.

**Step 3** Choose one of the following:

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

## Downgrade Your Software

This section describes how to downgrade.

- [About Activation Key Compatibility, page 35-25](#)
- [Perform the Downgrade, page 35-25](#)

## 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.

## Perform the Downgrade

To downgrade, perform the following steps.

### Procedure

- 
- |               |                                                                                                                                                                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Choose <b>Tools &gt; Downgrade Software</b> .<br>The Downgrade Software dialog box appears.                                                                                                                                                                                                                                                                |
| <b>Step 2</b> | For the ASA Image, click <b>Select Image File</b> .<br>The Browse File Locations dialog box appears.                                                                                                                                                                                                                                                       |
| <b>Step 3</b> | Click one of the following radio buttons: <ul style="list-style-type: none"><li>• <b>Remote Server</b>—Choose <b>ftp</b>, <b>smb</b>, or <b>http</b> from the drop-down list, and type the path to the old image file.</li><li>• <b>Flash File System</b>—Click <b>Browse Flash</b> to choose the old image file on the local flash file system.</li></ul> |
| <b>Step 4</b> | For the Configuration, click <b>Browse Flash</b> to choose the pre-migration configuration file. (By default this was saved on disk0).                                                                                                                                                                                                                     |
| <b>Step 5</b> | (Optional) In the Activation Key field, enter the old activation key if you need to revert to a pre-8.3 activation key.<br>See <a href="#">About Activation Key Compatibility, page 35-25</a> for more information.                                                                                                                                        |
| <b>Step 6</b> | Click <b>Downgrade</b> .<br>This tool is a shortcut for completing the following functions: <ol style="list-style-type: none"><li>1. Clearing the boot image configuration (<b>clear configure boot</b>).</li><li>2. Setting the boot image to be the old image (<b>boot system</b>).</li></ol>                                                            |

3. (Optional) Entering a new activation key (**activation-key**).
  4. Saving the running configuration to startup (**write memory**). This sets the BOOT environment variable to the old image, so when you reload, the old image is loaded.
  5. Copying the old configuration to the startup configuration (**copy old\_config\_url startup-config**).
  6. Reloading (**reload**).
- 

## Configure 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.

- [About Auto Update, page 35-26](#)
- [Guidelines for Auto Update, page 35-28](#)
- [Configure Communication with an Auto Update Server, page 35-28](#)

## About Auto Update

This section describes how Auto Update is implemented and why you might want to use Auto Update.

- [Auto Update Client or Server, page 35-26](#)
- [Auto Update Benefits, page 35-26](#)
- [Auto Update Server Support in Failover Configurations, page 35-27](#)

## Auto Update Client or Server

The ASA can be configured as either a client or a server. As an Auto Update client, it periodically polls the Auto Update Server for updates to software images and configuration files. As an Auto Update Server, it issues updates for ASAs configured as Auto Update clients.

## Auto Update Benefits

Auto Update is useful in solving many issues facing administrators for ASA management, such as:

- Overcoming dynamic addressing and NAT challenges.
- Committing configuration changes in one action.
- Providing a reliable method for updating software.
- Leveraging well-understood methods for high availability (failover).
- Providing flexibility with an open interface.
- Simplifying security solutions for Service Provider environments.

The Auto Update specification provides the infrastructure necessary for remote management applications to download ASA configurations, software images, and to perform basic monitoring from a centralized location or multiple locations.

The Auto Update specification allows the Auto Update server to either push configuration information and send requests for information to the ASA, or to pull configuration information by having the ASA periodically poll the Auto Update server. The Auto Update server can also send a command to the ASA to send an immediate polling request at any time. Communication between the Auto Update server and the ASA requires a communications path and local CLI configuration on each ASA.

## Auto Update Server Support in Failover Configurations

You can use the Auto Update Server to deploy software images and configuration files to ASAs in an Active/Standby failover configuration. To enable Auto Update on an Active/Standby failover configuration, enter the Auto Update Server configuration on the primary unit in the failover pair.

The following restrictions and behaviors apply to Auto Update Server support in failover configurations:

- Only single mode, Active/Standby configurations are supported.
- When loading a new platform software image, the failover pair stops passing traffic.
- When using LAN-based failover, new configurations must not change the failover link configuration. If they do, communication between the units will fail.
- Only the primary unit will perform the call home to the Auto Update Server. The primary unit must be in the active state to call home. If it is not, the ASA automatically fails over to the primary unit.
- Only the primary unit downloads the software image or configuration file. The software image or configuration is then copied to the secondary unit.
- The interface MAC address and hardware-serial ID is from the primary unit.
- The configuration file stored on the Auto Update Server or HTTP server is for the primary unit only.

## Auto Update Process Overview

The following is an overview of the Auto Update process in failover configurations. This process assumes that failover is enabled and operational. The Auto Update process cannot occur if the units are synchronizing configurations, if the standby unit is in the failed state for any reason other than SSM card failure, or if the failover link is down.

1. Both units exchange the platform and ASDM software checksum and version information.
2. The primary unit contacts the Auto Update Server. If the primary unit is not in the active state, the ASA first fails over to the primary unit and then contacts the Auto Update Server.
3. The Auto Update Server replies with software checksum and URL information.
4. If the primary unit determines that the platform image file needs to be updated for either the active or standby unit, the following occurs:
  - a. The primary unit retrieves the appropriate files from the HTTP server using the URL from the Auto Update Server.
  - b. The primary unit copies the image to the standby unit and then updates the image on itself.
  - c. If both units have new image, the secondary (standby) unit is reloaded first.
    - If hitless upgrade can be performed when secondary unit boots, then the secondary unit becomes the active unit and the primary unit reloads. The primary unit becomes the active unit when it has finished loading.
    - If hitless upgrade cannot be performed when the standby unit boots, then both units reload at the same time.

- d. If only the secondary (standby) unit has new image, then only the secondary unit reloads. The primary unit waits until the secondary unit finishes reloading.
  - e. If only the primary (active) unit has new image, the secondary unit becomes the active unit, and the primary unit reloads.
  - f. The update process starts again at Step 1.
5. If the ASA determines that the ASDM file needs to be updated for either the primary or secondary unit, the following occurs:
  - a. The primary unit retrieves the ASDM image file from the HTTP server using the URL provided by the Auto Update Server.
  - b. The primary unit copies the ASDM image to the standby unit, if needed.
  - c. The primary unit updates the ASDM image on itself.
  - d. The update process starts again at Step 1.
6. If the primary unit determines that the configuration needs to be updated, the following occurs:
  - a. The primary unit retrieves the configuration file from the using the specified URL.
  - b. The new configuration replaces the old configuration on both units simultaneously.
  - c. The update process begins again at Step 1.
7. If the checksums match for all image and configuration files, no updates are required. The process ends until the next poll time.

## Guidelines for Auto Update

- If the ASA configuration is updated from an Auto Update server, ASDM is not notified. You must choose **Refresh** or **File > Refresh ASDM with the Running Configuration on the Device** to obtain the latest configuration, and any changes to the configuration made in ASDM will be lost.
- If HTTPS is chosen as the protocol to communicate with the Auto Update server, the ASA uses SSL, which requires the ASA to have a DES or 3DES license.
- Auto Update is supported in single context mode only.

## Configure Communication with an Auto Update Server

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > System Image/Configuration > Auto Update**.
- The Auto Update pane consists of an Auto Update Servers table and two areas: the Timeout area and the Polling area.
- The Auto Update Servers table lets you view the parameters of previously configured Auto Update servers. The ASA polls the server listed at the top of the table first.
- Step 2** To change the order of the servers in the table, click **Move Up** or **Move Down**.
- The Auto Update Servers table includes the following columns:
- Server—The name or IP address of the Auto Update server.
  - User Name—The user name used to access the Auto Update server.



- Interface—The interface used when sending requests to the Auto Update server.
  - Verify Certificate—Indicates whether the ASA checks the certificate returned by the Auto Update server with the CA root certificates. The Auto Update server and the ASA must use the same CA.
- Step 3** Double-clicking any of the rows in the Auto Update Server table opens the Edit Auto Update Server dialog box, in which you can modify the Auto Update server parameters. These changes are immediately reflected in the table, but you must click **Apply** to save them to the configuration.
- Step 4** The Timeout area lets you set the amount of time the ASA waits for the Auto Update server to time out. The Timeout area includes the following fields:
- Enable Timeout Period—Check to enable the ASA to time out if no response is received from the Auto Update server.
  - Timeout Period (Minutes)—Enter the number of minutes the ASA will wait to time out if no response is received from the Auto Update server.
- Step 5** The Polling area lets you configure how often the ASA will poll for information from the Auto Update server. The Polling area includes the following fields:
- Polling Period (minutes)—The number of minutes the ASA will wait to poll the Auto Update server for new information.
  - Poll on Specified Days—Allows you to specify a polling schedule.
  - Set Polling Schedule—Displays the Set Polling Schedule dialog box where you can configure the days and time-of-day to poll the Auto Update server.
  - Retry Period (minutes)—The number of minutes the ASA will wait to poll the Auto Update server for new information if the attempt to poll the server fails.
  - Retry Count—The number of times the ASA will attempt to retry to poll the Auto Update server for new information.
- Step 6** Setting the Polling Schedule
- The Set Polling Schedule dialog box lets you configure specific days and the time-of-day for the ASA to poll the Auto Update server.
- The Set Polling Schedule dialog box includes the following fields:
- Days of the Week—Check the days of the week that you want the ASA to poll the Auto Update server.
- The Daily Update pane group lets you configure the time of day when you want the ASA to poll the Auto Update server, and includes the following fields:
- Start Time—Enter the hour and minute to begin the Auto Update poll.
  - Enable randomization—Check to enable the ASA to randomly choose a time to poll the Auto Update server.
- 
- 

## Monitoring Auto Update

You can use the **debug auto-update client** or **debug fover cmd-exe** commands to display the actions performed during the Auto Update process. The following is sample output from the **debug auto-update client** command. Run **debug** commands from a terminal session.

```
Auto-update client: Sent DeviceDetails to /cgi-bin/dda.pl of server 192.168.0.21
```

```

Auto-update client: Processing UpdateInfo from server 192.168.0.21
 Component: asdm, URL: http://192.168.0.21/asdm.bint, checksum:
0x94bced0261cc992ae710faf8d244cf32
 Component: config, URL: http://192.168.0.21/config-rms.xml, checksum:
0x67358553572688a805a155af312f6898
 Component: image, URL: http://192.168.0.21/cdisk73.bin, checksum:
0x6d091b43ce96243e29a62f2330139419
Auto-update client: need to update img, act: yes, stby yes
name
ciscoasa(config)# Auto-update client: update img on stby unit...
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 1501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 2001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 2501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 3001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 3501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 4001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 4501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 5001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 5501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 6001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 6501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 7001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 7501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 8001, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 8501, len = 1024
auto-update: Fover copyfile, seq = 4 type = 1, pseq = 9001, len = 1024
auto-update: Fover file copy waiting at clock tick 6129280
fover_parse: Rcvd file copy ack, ret = 0, seq = 4
auto-update: Fover filecopy returns value: 0 at clock tick 6150260, upd time 145980 msecs
Auto-update client: update img on active unit...
fover_parse: Rcvd image info from mate
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
auto-update: HA safe reload: reload active waiting with mate state: 20
Beginning configuration replication: Sending to mate.
auto-update: HA safe reload: reload active waiting with mate state: 50
auto-update: HA safe reload: reload active waiting with mate state: 50

auto-update: HA safe reload: reload active waiting with mate state: 80
 Sauto-update: HA safe reload: reload active unit at clock tick: 6266860
Auto-update client: Succeeded: Image, version: 0x6d091b43ce96243e29a62f2330139419

```

The following syslog message is generated if the Auto Update process fails:

```
%ASA4-612002: Auto Update failed: file version: version reason: reason
```

The *file* is “image”, “asdm”, or “configuration”, depending on which update failed. The *version* is the version number of the update. And the *reason* is the reason that the update failed.

# History for Software and Configurations

**Table 35-1** History for Software and Configurations

| Feature Name                                                   | Platform Releases | Feature Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Secure Copy client                                             | 9.1(5)/9.2(1)     | <p>The ASA now supports the Secure Copy (SCP) client to transfer files to and from a SCP server.</p> <p>We modified the following screens:</p> <p>Tools &gt; File Management &gt; File Transfer &gt; Between Remote Server and Flash</p> <p><b>Configuration &gt; Device Management &gt; Management Access &gt; File Access &gt; Secure Copy (SCP) Server</b></p>                                                                                                                                                                                                                                                                                                                                                                                                      |
| Auto Update server certificate verification enabled by default | 9.2(1)            | <p>The Auto Update server certificate verification is now enabled by default; for new configurations, you must explicitly disable certificate verification. If you are upgrading from an earlier release, and you did not enable certificate verification, then certificate verification is not enabled, and you see the following warning:</p> <p>WARNING: The certificate provided by the auto-update servers will not be verified. In order to verify this certificate please use the verify-certificate option.</p> <p>The configuration will be migrated to explicitly configure no verification.</p> <p>We modified the following screen: Configuration &gt; Device Management &gt; System/Image Configuration &gt; Auto Update &gt; Add Auto Update Server.</p> |
| System backup and restore using the CLI                        | 9.3(2)            | <p>You can now back up and restore complete system configurations, including images and certificates, using the CLI.</p> <p>We did not modify any ASDM screens.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Recovering the ASA 5506W configuration                         | 9.4(1)            | <p>We now support recovery of the ASA 5506W configuration.</p> <p>We did not modify any ASDM screens.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Recovering and loading a new ASA 5506W image                   |                   | <p>We now support the recovery and loading of a new ASA 5506W image.</p> <p>We did not modify any ASDM screens.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |





# Response Automation for System Events

This chapter describes how to configure the Embedded Event Manager (EEM).

- [About the EEM, page 36-1](#)
- [Guidelines for the EEM, page 36-2](#)
- [Configure the EEM, page 36-3](#)
- [Monitoring the EEM, page 36-5](#)
- [History for the EEM, page 36-6](#)

## About the EEM

The EEM service enables you to debug problems and provides general purpose logging for troubleshooting. There are two components: events to which the EEM responds or listens, and event manager applets that define actions as well as the events to which the EEM responds. You may configure multiple event manager applets to respond to different events and perform different actions.

## Supported Events

The EEM supports the following events:

- Syslog—The ASA uses syslog message IDs to identify syslog messages that trigger an event manager applet. You may configure multiple syslog events, but the syslog message IDs may not overlap within a single event manager applet.
- Timers—You may use timers to trigger events. You may configure each timer only once for each event manager applet. Each event manager applet may have up to three timers. The three types of timers are the following:
  - Watchdog (periodic) timers trigger an event manager applet after the specified time period following the completion of the applet actions and restart automatically.
  - Countdown (one-shot) timers trigger an event manager applet once after the specified time period and do not restart unless they are removed, then re-added.
  - Absolute (once-a-day) timers cause an event to occur once a day at a specified time, and restart automatically. The time-of-day format is in hh:mm:ss.

You may configure only one timer event of each type for each event manager applet.

- None—The none event is triggered when you run an event manager applet manually using the CLI or ASDM.

- **Crash**—The crash event is triggered when the ASA crashes. Regardless of the value of the **output** command, the **action** commands are directed to the crashinfo file. The output is generated before the **show tech** command.

## Actions on Event Manager Applets

When an event manager applet is triggered, the actions on the event manager applet are performed. Each action has a number that is used to specify the sequence of the actions. The sequence number must be unique within an event manager applet. You may configure multiple actions for an event manager applet. The commands are typical CLI commands, such as **show blocks**.

## Output Destinations

You may send the output from the actions to a specified location using the **output** command. Only one output value may be enabled at any one time. The default value is **output none**. This value discards any output from the **action** commands. The command runs in global configuration mode as a user with privilege level 15 (the highest). The command may not accept any input, because it is disabled. You may send the output of the **action** CLI commands to one of three locations:

- **None**, which is the default and discards the output
- **Console**, which sends the output to the ASA console
- **File**, which sends the output to a file. The following four file options are available:
  - **Create a unique file**, which creates a new, uniquely named file each time that an event manager applet is invoked
  - **Create/overwrite a file**, which overwrites a specified file each time that an event manager applet is invoked.
  - **Create/append to a file**, which appends to a specified file each time that an event manager applet is invoked. If the file does not yet exist, it is created.
  - **Create a set of files**, which creates a set of uniquely named files that are rotated each time that an event manager applet is invoked.

## Guidelines for the EEM

This section describes guidelines and limitations that you should check before configuring the EEM.

### Context Mode Guidelines

Not supported in multiple context mode.

### Additional Guidelines

- During a crash, the state of the ASA is generally unknown. Some commands may not be safe to run during this condition.
- The name of an event manager applet may not contain spaces.
- You cannot modify the None event and Crashinfo event parameters.
- Performance may be affected because syslog messages are sent to the EEM for processing.
- The default output is **output none** for each event manager applet. To change this setting, you must enter a different output value.

- You may have only one output option defined for each event manager applet.

## Configure the EEM

Configuring the EEM consists of the following tasks:

- 
- |               |                                                                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Create an event manager applet, then configure various events. See <a href="#">Create an Event Manager Applet and Configure Events</a> , page 36-3.                                                     |
| <b>Step 2</b> | Configure an action on an event manager applet, then configure a destination for output from an action. See <a href="#">Configure an Action and Destinations for Output from an Action</a> , page 36-4. |
| <b>Step 3</b> | Run an event manager applet. See <a href="#">Run an Event Manager Applet</a> , page 36-5.                                                                                                               |
| <b>Step 4</b> | Track memory allocations and memory usage for the EEM. See <a href="#">Track Memory Allocation and Memory Usage</a> , page 36-5.                                                                        |
- 

## Create an Event Manager Applet and Configure Events

To create an event manager applet and configure events, perform the following steps:

### Procedure

- 
- |               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | In ASDM, choose <b>Configuration &gt; Device Management &gt; Advanced &gt; Embedded Event Manager</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Step 2</b> | Click <b>Add</b> to display the <b>Add Event Manager Applet</b> dialog box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 3</b> | Enter the name of the applet (without spaces) and describe what it does. The description may be up to 256 characters long. You may include spaces in description text if it is placed within quotes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Step 4</b> | Click <b>Add</b> in the <b>Events</b> area to display the <b>Add Event Manager Applet Event</b> dialog box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 5</b> | Choose the event type that you want to configure from the <b>Type</b> drop-down list. The available options are <b>crashinfo</b> , <b>None</b> , <b>Syslog</b> , <b>Once-a-day timer</b> , <b>One-shot timer</b> , and <b>Periodic</b> timer. <ul style="list-style-type: none"><li>• <b>Syslog</b>: Enter a single syslog message or a range of syslog messages. If a syslog message occurs that matches the specified individual syslog message or range of syslog messages, an event manager applet is triggered. (Optional) Enter the number of times in the <b>occurrences</b> field that the syslog message must occur for an event manager applet to be invoked. The default is 1 occurrence every 0 seconds. Valid values are from 1 - 4294967295. (Optional) Enter the number of seconds in the <b>period</b> field within which the syslog messages must occur to invoke the action. This value limits how frequently an event manager applet is invoked to at most once in the configured period. Valid values are from 0 - 604800. A value of 0 means that no period is defined.</li><li>• <b>Periodic</b>: Enter the time period in seconds. The number of seconds may range from 1- 604800.</li><li>• <b>Once-a-day timer</b>: Enter the time of day in hh:mm:ss. The time range is from 00:00:00 (midnight) to 23:59:59.</li><li>• <b>One-shot timer</b>: Enter the time period in seconds. The number of seconds may range from 1- 604800.</li><li>• <b>None</b>: Choose this option to invoke an event manager applet manually.</li></ul> |

- **crashinfo:** Choose this option to trigger a crash event when the ASA crashes.

## Configure an Action and Destinations for Output from an Action

To configure an action and specific destinations for sending output from an action, perform the following steps:

### Procedure

#### Step 1

**Step 1** Click **Add** to display the **Add Event Manager Applet** dialog box.

**Step 2** Enter the name of the applet (without spaces) and describe what it does. The description may be up to 256 characters long.

**Step 3** Click **Add** in the **Actions** area to display the **Add Event Manager Applet Action** dialog box.

**Step 4** Enter the unique sequence number in the **Sequence #** field. Valid sequence numbers range from 0 - 4294967295.

**Step 5** Enter the CLI command in the **CLI Command** field. The command runs in global configuration mode as a user with privilege level 15 (the highest). The command may not accept any input, because it is disabled.

**Step 6** Click **OK** to close the **Add Event Manager Applet Action** dialog box.

The newly added action appears in the **Actions** list.

**Step 7** Click **Add** to open the **Add Event Manager Applet** dialog box.

**Step 8** Choose one of the available output destination options:

- Choose the **None** option from the **Output Location** drop-down list to discard any output from the **action** commands. This is the default setting.
- Choose the **Console** option from the **Output Location** drop-down list to send the output of the **action** commands to the console.



**Note** Running this command affects performance.

- Choose the **File** option from the **Output Location** drop-down list to send the output of the **action** commands to a new file for each event manager applet that is invoked. The **Create a unique file** option is automatically selected as the default.

The filename has the format of *eem-applet-timestamp.log*, in which *applet* is the name of the event manager applet and *timestamp* is a dated time stamp in the format of YYYYMMDD-hhmmss.

- Choose the **File** option from the **Output Location** drop-down list, then choose the **Create a set of files** option from the drop-down list to create a set of rotated files.

When a new file is to be written, the oldest file is deleted, and all subsequent files are renumbered before the first file is written. The newest file is indicated by 0, and the oldest file is indicated by the highest number. Valid values for the rotate value range from 2 - 100. The filename format is *eem-applet-x.log*, in which *applet* is the name of the applet, and *x* is the file number.



- Choose the **File** option from the **Output Location** drop-down list, then choose the **Create/overwrite a file** option from the drop-down list to write the **action** command output to a single file, which is overwritten every time.
- Choose the **File** option from the **Output Location** drop-down list, then choose the **Create/append a file** option from the drop-down list to writes the **action** command output to a single file, but that file is appended to every time.

**Step 9** Click **OK** to close the **Add Event Manager Applet** dialog box.

The specified output destination appears in the **Embedded Event Manager** pane.

---

## Run an Event Manager Applet

To run an event manager applet, perform the following steps:

### Procedure

**Step 1** In the **Embedded Event Manager** pane, select an event manager applet from the list that has been configured with the **None** event.

**Step 2** Click **Run**.

---

## Track Memory Allocation and Memory Usage

To log memory allocation and memory usage, perform the following steps:

### Procedure

**Step 1** Choose **Configuration > Device Management > Advanced > Embedded Event Manager**.

**Step 2** Click **Add** to display the **Add Event Manager Applet** dialog box.

**Step 3** Click **Add** again to display the **Add Event Manager Applet Event** dialog box.

**Step 4** Choose **memory-logging-wrap** from the drop-down list.

**Step 5** Click **OK** to add it to the **Events** list.

**Step 6** Click **OK** again to add it to the **Applets** list.

---

## Monitoring the EEM

See the following screens to monitor the EEM.

- **Monitoring > Properties > EEM Applets**

This pane shows the list of EEM applets and their hit count value.

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

## History for the EEM

**Table 36-1** History for the EEM

| Feature Name                 | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Embedded Event Manager (EEM) | 9.2(1)            | <p>The EEM service enables you to debug problems and provides general purpose logging for troubleshooting. There are two components: events to which the EEM responds or listens, and event manager applets that define actions as well as the events to which the EEM responds. You may configure multiple event manager applets to respond to different events and perform different actions.</p> <p>We introduced the following screens: Configuration &gt; Device Management &gt; Advanced &gt; Embedded Event Manager, Monitoring &gt; Properties &gt; EEM Applets.</p> |
| Memory tracking for the EEM  | 9.4(1)            | <p>We have added a new debugging feature to log memory allocations and memory usage, and to respond to memory logging wrap events.</p> <p>We modified the following screen: Configuration &gt; Device Management &gt; Advanced &gt; Embedded Event Manager &gt; Add Event Manager Applet &gt; Add Event Manager Applet Event.</p>                                                                                                                                                                                                                                            |



## Testing and Troubleshooting

This chapter describes how to troubleshoot the Cisco ASA and test basic connectivity.

- [Configure and Run Captures with the Packet Capture Wizard, page 37-1](#)
- [vCPU Usage in the ASAv, page 37-5](#)
- [Test Your Configuration, page 37-7](#)
- [Monitoring Performance and System Resources, page 37-14](#)
- [Monitoring Connections, page 37-17](#)

### Configure and Run Captures with the Packet Capture Wizard

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



#### Note

This tool does not support clientless SSL VPN capture.

To configure and run captures, perform the following steps:

#### Procedure

**Step 1** Choose **Wizards > Packet Capture Wizard**.

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

- Selecting an ingress interface.
- Selecting an egress interface.
- Setting the buffer parameters.
- Running the captures.
- Saving the captures to your PC (optional).

**Step 2** Click **Next**.

In a clustering environment, the **Cluster Option** screen appears. Go to [Step 3](#).

In a non-clustering environment, the **Ingress Traffic Selector** screen appears. Go to [Step 4](#).

- Step 3** Choose one of the following options in the **Cluster Option** screen for running a capture: **This device only** or **The whole cluster**, then click **Next** to display the **Ingress Selector** screen.
- Step 4** Click the **Select Interface** radio button to capture packets on an interface. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.
- Step 5** Do one of the following in the **Packet Match Criteria** area:
- Click the **Specify access-list** radio button to specify the ACL to use for matching packets, then choose the ACL from the **Select ACL** drop-down list. Click **Manage** to display the **ACL Manager** pane to add a previously configured ACL to the current drop-down list. Choose an ACL, then click **OK**.
  - Click the **Specify Packet Parameters** radio button to specify packets parameters.
- Step 6** To continue, see [Ingress Traffic Selector, page 37-3](#).
- Step 7** Click **Next** to display the **Egress Traffic Selector** screen. To continue, see [Egress Traffic Selector, page 37-4](#).




---

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

---

- Step 8** Click **Next** to display the **Buffers & Captures** screen. To continue, see [Buffers, page 37-4](#).
- Step 9** Check the **Get capture every 10 seconds** check box in the **Capture Parameters** area to obtain the latest capture every 10 seconds automatically. By default, this capture uses the circular buffer.
- Step 10** You specify the buffer size and packet size in the **Buffer Parameters** area. The buffer size is the maximum amount of memory that the capture can use to store packets. The packet size is the longest packet that the capture can hold. We recommend that you use the longest packet size to capture as much information as possible.
- Enter the packet size. The valid size ranges from 14 - 1522 bytes.
  - Enter the buffer size. The valid size ranges from 1534 - 33554432 bytes.
  - Check the **Use circular buffer** check box to store captured packets.




---

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

---

- Step 11** Click **Next** to display the **Summary** screen, which shows the cluster options for all units in the cluster (if you are using clustering), traffic selectors, and buffer parameters that you have entered. To continue, see [Summary, page 37-5](#).
- Step 12** Click **Next** to display the **Run Captures** screen, then click **Start** to begin capturing packets. Click **Stop** to end the capture. To continue, see [Run Captures, page 37-5](#). If you are using clustering, go to Step 14.
- Step 13** Click **Get Capture Buffer** to determine how much buffer space you have remaining. Click **Clear Buffer on Device** to remove the current content and allow room in the buffer to capture more packets.
- Step 14** In a clustering environment, on the **Run Captures** screen, perform one or more of the following steps:
- Click **Get Cluster Capture Summary** to view a summary of packet capture information for all units in the cluster, followed by packet capture information for each unit.
  - Click **Get Capture Buffer** to determine how much buffer space you have remaining in each unit of the cluster. The **Capture Buffer from Device** dialog box appears.

- Click **Clear Capture Buffer** to remove the current content for one or all of the units in a cluster and allow room in the buffer to capture more packets.
- Step 15** Click **Save captures** to display the **Save Capture** dialog box. You have the option of saving either the ingress capture, the egress capture, or both. To continue, see [Save Captures, page 37-5](#).
- Step 16** Click **Save Ingress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.
- Step 17** Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools > Preferences** for analyzing the ingress capture.
- Step 18** Click **Save Egress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.
- Step 19** Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools > Preferences** for analyzing the egress capture.
- Step 20** Click **Close**, then click **Finish** to exit the wizard.
- 

## Ingress Traffic Selector

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

### Procedure

- 
- Step 1** Choose the ingress interface name from the drop-down list.
- Step 2** Enter the ingress source host and network. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.
- Step 3** Enter the ingress destination host and network.
- Step 4** Enter the protocol type to capture. Available protocols are ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.
- a. Enter the ICMP type for ICMP only. Available types include all, alternate address, conversion-error, echo, echo-reply, information-reply, information-request, mask-reply, mask-request, mobile-redirect, parameter-problem, redirect, router-advertisement, router-solicitation, source-quench, time-exceeded, timestamp-reply, timestamp-request, traceroute, or unreachable.
  - b. Specify the source and destination port services for the TCP and UDP protocols only. Available options include the following:
    - Choose **All Services** to include all services.
    - Choose **Service Groups** to include a service group.

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

- Step 5** Check the **SGT number** check box in the **Security Group Tagging** area and enter the security group tag number to enable packet capture for the Cisco TrustSec service. Valid security group tag numbers range from 2 - 65519.
- 

## Egress Traffic Selector

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

### Procedure

- 
- Step 1** Click the **Select Interface** radio button to capture packets on an interface. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.
- Step 2** Choose the egress interface name from the drop-down list.
- Step 3** Enter the egress source host and network.
- Step 4** Enter the egress destination host and network.

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

---

## Buffers

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

### Procedure

- 
- Step 1** Enter the longest packet that the capture can hold. Use the longest size available to capture as much information as possible.
- Step 2** Enter the maximum amount of memory that the capture can use to store packets.
- Step 3** Use the circular buffer to store packets. When the circular buffer has used all of the buffer storage, the capture will overwrite the oldest packets first.
-

## Summary

The **Summary** screen shows the cluster options (if you are using clustering), traffic selectors, and the buffer parameters for the packet capture selected in the previous wizard screens.

## Run Captures

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

### Procedure

- 
- |               |                                                                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>Start</b> to begin the packet capture session on a selected interface.                                                                                                               |
| <b>Step 2</b> | Click <b>Stop</b> to stop the packet capture session on a selected interface.                                                                                                                 |
| <b>Step 3</b> | Click <b>Get Capture Buffer</b> to obtain a snapshot of the captured packets on the interface.                                                                                                |
| <b>Step 4</b> | Click <b>Ingress</b> to show the capture buffer on the ingress interface.                                                                                                                     |
| <b>Step 5</b> | Click <b>Egress</b> to show the capture buffer on the egress interface.                                                                                                                       |
| <b>Step 6</b> | Click <b>Clear Buffer on Device</b> to clear the buffer on the device.                                                                                                                        |
| <b>Step 7</b> | Click <b>Launch Network Sniffer Application</b> to start the packet analysis application for analyzing the ingress capture or the egress capture specified in <b>Tools &gt; Preferences</b> . |
| <b>Step 8</b> | Click <b>Save Captures</b> to save the ingress and egress captures in either ASCII or PCAP format.                                                                                            |
- 

## Save Captures

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

### Procedure

- 
- |               |                                                                                                  |
|---------------|--------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Click <b>ASCII</b> to save the capture buffer in ASCII format.                                   |
| <b>Step 2</b> | Click <b>PCAP</b> to save the capture buffer in PCAP format.                                     |
| <b>Step 3</b> | Click <b>Save ingress capture</b> to specify a file in which to save the ingress packet capture. |
| <b>Step 4</b> | Click <b>Save egress capture</b> to specify a file in which to save the egress packet capture.   |
- 

## vCPU Usage in the ASAv

The ASAv vCPU usage shows the amount of vCPUs used for the data path, control point, and external processes.

The vSphere reported vCPU usage includes the ASAv usage as described plus:

- ASAv idle time
- %SYS overhead used for the ASAv VM
- Overhead of moving packets between vSwitches, vNICs, and pNICs. This overhead can be quite significant.

## CPU Usage Example

The following is an example in which the reported vCPU usage is substantially different:

- ASAv reports: 40%
- DP: 35%
- External Processes: 5%
- vSphere reports: 95%
- ASA (as ASAv reports): 40%
- ASA idle polling: 10%
- Overhead: 45%

The overhead is used to perform hypervisor functions and to move packets between NICs and vNICs using the vSwitch.

Usage can exceed 100% because the ESXi server can use additional compute resources for overhead on behalf of the ASAv.

## VMware CPU Usage Reporting

In vSphere, click the **VM Performance** tab, then click **Advanced** to display the **Chart Options** drop-down list, which shows vCPU usage for each state (%USER, %IDLE, %SYS, and so on) of the VM. This information is useful for understanding VMware's perspective on where CPU resources are being used.

On the ESXi server shell (you access the shell by using SSH to connect to the host), `esxtop` is available. `Esxtop` has a similar look and feel to the Linux `top` command and provides VM state information for vSphere performance, including the following:

- Details on vCPU, memory, and network usage
- vCPU usage for each state of each VM.
- Memory (type M while running) and network (type N while running), as well as statistics and the number of RX drops

## ASAv and vCenter Graphs

There are differences in the CPU % numbers between the ASAv and vCenter:

- The vCenter graph numbers are always higher than the ASAv numbers.
- vCenter calls it %CPU usage; the ASAv calls it %CPU utilization.

The terms “%CPU utilization” and “%CPU usage” mean different things:

- CPU utilization provides statistics for physical CPUs.



- CPU usage provides statistics for logical CPUs, which is based on CPU hyperthreading. But because only one vCPU is used, hyperthreading is not turned on.

vCenter calculates the CPU % usage as follows:

Amount of actively used virtual CPUs, specified as a percentage of the total available CPUs

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is as follows:

Usage in MHz / number of virtual CPUs x core frequency

When you compare the usage in MHz, both the vCenter and ASA values match. According to the vCenter graph, MHz % CPU usage is calculated as:

$$60 / (2499 \times 1 \text{ vCPU}) = 2.4$$

## Test Your Configuration

This section describes how to test connectivity for the single mode ASA or for each security context, how to ping the ASA interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

- [Test Basic Connectivity: Pinging Addresses, page 37-7](#)
- [Trace Routes to Hosts, page 37-12](#)
- [Tracing Packets to Test Policy Configuration, page 37-14](#)

## Test Basic Connectivity: Pinging Addresses

Ping is a simple command that lets you determine if a particular address is alive and responsive. The following topics explain more about the command and what types of testing you can accomplish with it.

- [What You Can Test Using Ping, page 37-7](#)
- [Choosing Between ICMP and TCP Ping, page 37-8](#)
- [Enable ICMP, page 37-8](#)
- [Ping Hosts, page 37-9](#)
- [Test ASA Connectivity Systematically, page 37-10](#)

## What You Can Test Using Ping

When you ping a device, a packet is sent to the device and the device returns a reply. This process enables network devices to discover, identify, and test each other.

You can use ping to do the following tests:

- Loopback testing of two interfaces—You can initiate a ping from one interface to another on the same ASA, as an external loopback test to verify basic “up” status and operation of each interface.
- Pinging to an ASA—You can ping an interface on another ASA to verify that it is up and responding.

- Pinging through an ASA—You can ping through an intermediate ASA by pinging a device on the other side of the ASA. The packets will pass through two of the intermediate ASA's interfaces as they go in each direction. This action performs a basic test of the interfaces, operation, and response time of the intermediate unit.
- Pinging to test questionable operation of a network device—You can ping from an ASA interface to a network device that you suspect is functioning incorrectly. If the interface is configured correctly and an echo is not received, there might be problems with the device.
- Pinging to test intermediate communications—You can ping from an ASA interface to a network device that is known to be functioning correctly. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

## Choosing Between ICMP and TCP Ping

The ASA includes the traditional ping, which sends ICMP Echo Request packets and gets Echo Reply packets in return. This is the standard tool and works well if all intervening network devices allow ICMP traffic. With ICMP ping, you can ping IPv4 or IPv6 addresses, or host names.

However, some networks prohibit ICMP. If this is true of your network, you can instead use TCP ping to test network connectivity. With TCP ping, the ping sends TCP SYN packets, and considers the ping a success if it receives a SYN-ACK in response. With TCP ping, you can ping IPv4 addresses or host names, but you cannot ping IPv6 addresses.

Keep in mind that a successful ICMP or TCP ping simply means that the address you are using is alive and responding to that specific type of traffic. This means that basic connectivity is working. Other policies running on a device could prevent specific types of traffic from successfully getting through a device.

## Enable ICMP

By default, you can ping from a high security interface to a low security interface. You just need to enable ICMP inspection to allow returning traffic through. If you want to ping from low to high, then you need to apply an ACL to allow traffic.

When pinging an ASA interface, any ICMP rules applied to the interface must allow Echo Request and Echo Response packets. ICMP rules are optional: if you do not configure them, all ICMP traffic to an interface is allowed.

This procedure explains all of ICMP configuration you might need to complete to enable ICMP pinging of ASA interfaces, or for pinging through an ASA.

### Procedure

- 
- Step 1** Ensure ICMP rules allow Echo Request/Echo Response.
- ICMP rules are optional and apply to ICMP packets sent directly to an interface. If you do not apply ICMP rules, all ICMP access is allowed. In this case, no action is required.
- However, if you do implement ICMP rules, ensure that you include rules that permit any address for the Echo and Echo-Reply messages on each interface. Configure ICMP rules on the **Configuration > Device Management > Management Access > ICMP** page.
- Step 2** Ensure access rules allow ICMP.

When pingging a host through an ASA, access rules must allow ICMP traffic to leave and return. The access rule must at least allow Echo Request/Echo Reply ICMP packets. You can add these rules as global rules.

If you do not have access rules, you will need to also allow the other type of traffic you want, because applying any access rules to an interface adds an implicit deny, so all other traffic will be dropped.

Configure access rules on the **Configuration > Firewall > Access Rules** page. If you are simply adding the rules for testing purposes, you can delete them after completing the tests.

**Step 3** Enable ICMP inspection.

ICMP inspection is needed when pingging through the ASA, as opposed to pingging an interface. Inspection allows returning traffic (that is, the Echo Reply packet) to return to the host that initiated the ping, and also ensures there is one response per packet, which prevents certain types of attack.

You can simply enable ICMP inspection in the default global inspection policy.

- a. Choose **Configuration > Firewall > Service Policy Rules**.
  - b. Edit the **inspection\_default** global rule.
  - c. On the **Rule Actions > Protocol Inspection** tab, select ICMP.
  - d. Click **OK**, then **Apply**.
- 

## Ping Hosts

To ping any device, you simply choose **Tools > Ping**, enter the IP address or host name of the destination you are pingging, and click **Ping**. For TCP ping, you select **TCP** and also include the destination port. That is usually the extent of any test you need to run.

Example output for a successful ping:

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

If the ping fails, the output indicates ? for each failed attempt, and the success rate is less than 100 percent (complete failure is 0 percent):

```
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
????
Success rate is 0 percent (0/5)
```

However, you can also add parameters to control some aspects of the ping. Following are your basic options:

- ICMP ping—You can select the interface through which the destination host is connected. If you do not select an interface, the routing table is used to determine the correct interface. You can ping IPv4 or IPv6 addresses or host names.
- TCP ping—You must also select the TCP port for the destination you are pingging. For example, **www.example.com 80** to ping the HTTP port. You can ping IPv4 addresses or host names, but not IPv6 addresses.

You also have the option to specify the source address and port that is sending the ping. In this case, optionally select the interface through which the source sends the ping (the routing table is used when you do not select an interface).

Finally, you can specify how often to repeat the ping (the default is 5 times) or the timeout for each attempt (the default is 2 seconds).

## Test ASA Connectivity Systematically

If you want to do a more systematic test of ASA connectivity, you can use the following general procedure.


### Before You Begin

If you want to see the syslog messages mentioned in the procedure, enable logging (the **logging enable** command, or **Configuration > Device Management > Logging > Logging Setup** in ASDM).

### Procedure


- Step 1** Draw a diagram of your single-mode ASA or security context that shows the interface names, security levels, and IP addresses. The diagram should also include any directly connected routers and a host on the other side of the router from which you will ping the ASA.

**Figure 37-1 Network Diagram with Interfaces, Routers, and Hosts**




- Step 2** Ping each ASA interface from the directly connected routers. For transparent mode, ping the management IP address. This test ensures that the ASA interfaces are active and that the interface configuration is correct.

A ping might fail if the ASA interface is not active, the interface configuration is incorrect, or if a switch between the ASA and a router is down (see the following figure). In this case, no debugging messages or syslog messages appear, because the packet never reaches the ASA.

**Figure 37-2 Ping Failure at the ASA Interface**

If the ping reply does not return to the router, then a switch loop or redundant IP addresses might exist (see the following figure).

**Figure 37-3 Ping Failure Because of IP Addressing Problems**

- Step 3** Ping each ASA interface from a remote host. For transparent mode, ping the management IP address. This test checks whether the directly connected router can route the packet between the host and the ASA, and whether the ASA can correctly route the packet back to the host.

A ping might fail if the ASA does not have a return route to the host through the intermediate router (see the following figure). In this case, the debugging messages show that the ping was successful, but syslog message 110001 appears, indicating a routing failure has occurred.

**Figure 37-4 Ping Failure Because the ASA Has No Return Route**

- Step 4** Ping from an ASA interface to a network device that you know is functioning correctly.
- If the ping is not received, a problem with the transmitting hardware or interface configuration may exist.
  - If the ASA interface is configured correctly and it does not receive an echo reply from the “known good” device, problems with the interface hardware receiving function may exist. If a different interface with “known good” receiving capability can receive an echo after pinging the same “known good” device, the hardware receiving problem of the first interface is confirmed.
- Step 5** Ping from the host or router through the source interface to another host or router on another interface. Repeat this step for as many interface pairs as you want to check. If you use NAT, this test shows that NAT is operating correctly.

If the ping succeeds, a syslog message appears to confirm the address translation for routed mode (305009 or 305011) and that an ICMP connection was established (302020). You can also enter either the **show xlate** or **show conns** command to view this information.

The ping might fail because NAT is not configured correctly. In this case, a syslog message appears, showing that the NAT failed (305005 or 305006). If the ping is from an outside host to an inside host, and you do not have a static translation, you get message 106010.

**Figure 37-5 Ping Failure Because the ASA is Not Translating Addresses**



## Trace Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path.

### Procedure

- 
- Step 1** [Make the ASA Visible on Trace Routes, page 37-12.](#)
- Step 2** [Determine Packet Routes, page 37-13.](#)
- 

## Make the ASA Visible on Trace Routes

By default, the ASA does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the ASA, and increase the rate limit on ICMP unreachable messages.

### Procedure

- 
- Step 1** Decrement the TTL using a service policy.
- Choose **Configuration > Firewall > Service Policy Rules**.
  - Add or edit a rule. For example, if you already have a rule to which you can add the option to decrement TTL, you do not need to create a new one.
  - Progress through the wizard to the Rule Actions page, applying the rule globally or to an interface, and specifying the traffic match. For example, you could create a global match any rule.
  - On the Rule Actions page, click the **Connection Settings** tab, and select **Decrement time to live for a connection**.
  - Click **OK** or **Finish**, then **Apply**.
- Step 2** Increase the ICMP unreachable rate limit.
- Choose **Configuration > Device Management > Management Access > ICMP**.
  - Increase the **IPv4 ICMP Unreachable Message Limits > Rate Limit** value at the bottom of the page. For example, increase it to 50.

- c. Click **Apply**.

## Determine Packet Routes

Use Traceroute to help you to determine the route that packets will take to their destination. A traceroute works by sending UDP packets to a destination on an invalid port. Because the port is not valid, the routers along the way to the destination respond with an ICMP Time Exceeded Message, and report that error to the ASA.

The traceroute shows the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following table explains the output symbols.

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

### Procedure

- Step 1** Choose **Tools > Traceroute**.
- Step 2** Enter the destination hostname or IP address to which you are tracing the route. Configure a DNS server to use a host name.
- Step 3** (Optional) Configure the characteristics of the trace. The defaults are appropriate in most cases.
- **Timeout**—How long to wait for a response before timing out. The default is 3 seconds.
  - **Port**—The UDP port to use. The default is 33434.
  - **Probe**—How many probes to send at each TTL level. The default is 3.
  - **TTL**—The minimum and maximum time-to-live values for the probes. The minimum default is one, but you can set it to a higher value to suppress the display of known hops. The maximum default is 30. The traceroute terminates when the packet reaches the destination or when the maximum value is reached.
  - **Specify source interface or IP address**—The interface to use as the source of the trace. You can specify the interface by name or by IP address. In transparent mode, you must use the management address.
  - **Reverse Resolve**—Whether to have the output display the names of hops encountered if DNS name resolution is configured. Deselect the option to show IP addresses only.
  - **Use ICMP**—Whether to send ICMP probe packets instead of UDP probe packets.
- Step 4** Click **Trace Route** to start the traceroute.

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

---

## Tracing Packets to Test Policy Configuration

You can test your policy configuration by modeling a packet based on source and destination addressing and protocol characteristics. The trace does policy lookup to test access rules, NAT, routing, and so forth, to see if the packet would be permitted or denied.

By testing packets this way, you can see the results of your policies and test whether the types of traffic you want to allow or deny are handled as desired. Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied when they should be allowed.

### Procedure

---

- Step 1** Choose **Tools > Packet Tracer**.
- Step 2** Choose the source interface for the packet trace.
- Step 3** Specify the protocol type for the packet trace. Available protocol types include ICMP, IP, TCP, and UDP.
- Step 4** (Optional.) If you want to trace a packet where the security group tag value is embedded in the Layer 2 CMD header (Trustsec), check **SGT number** and enter the security group tag number, 0-65533.
- Step 5** Specify the source and destination for the packets.  
You can specify IPv4 or IPv6 addresses, fully-qualified domain names (FQDN), or security group names or tags, if you use Cisco Trustsec. For the source address, you can also specify a username in the format Domain\username.
- Step 6** Specify the protocol characteristics:
- ICMP—Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
  - TCP/UDP—Enter the source and destination port numbers.
  - Raw IP—Enter the protocol number, 0-255.
- Step 7** Click **Start** to trace the packet.
- The Information Display Area shows detailed messages about the results of the packet trace.
- 

## Monitoring Performance and System Resources

You can monitor a variety of system resources to identify performance or other potential problems.



## Monitoring Performance

You can view ASA performance information in a graphical or tabular format.

### Procedure

- 
- Step 1** Choose **Monitoring > Properties > Connection Graphs > Perfmon**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Select up to four entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- AAA Perfmon—Requests per second for authentication, authorization, and accounting requests.
  - Inspection Perfmon—Packets per second for HTTP, FTP, and TCP inspection.
  - Web Perfmon—Requests per second for URL access and URL server requests.
  - Connections Perfmon—Connections per second for all connections, UDP connections, TCP connections, and TCP Intercept.
  - Xlate Perfmon—NAT xlates per second.
- Step 4** Click **Show Graphs**.
- You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
- 

## Monitoring Memory Blocks

You can view free and used memory blocks information in a graphical or tabular format.

### Procedure

- 
- Step 1** Choose **Monitoring > Properties > System Resources Graphs > Blocks**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- Blocks Used—Displays the ASA used memory blocks.
  - Blocks Free—Displays the ASA free memory blocks.
- Step 4** Click **Show Graphs**.
- You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
-

## Monitoring CPU

You can view CPU utilization.

### Procedure

- 
- Step 1** Choose **Monitoring > Properties > System Resources Graphs > CPU**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Add CPU Utilization to the Selected Graphs list.
- Step 4** Click **Show Graphs**.
- You can toggle the graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
- 

## Monitoring Memory

You can view memory utilization information in a graphical or tabular format.

### Procedure

- 
- Step 1** Choose **Monitoring > Properties > System Resources Graphs > Memory**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- Free Memory—Displays the ASA free memory.
  - Used Memory—Displays the ASA used memory.
- Step 4** Click **Show Graphs**.
- You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
-

## Monitoring Per-Process CPU Usage

You can monitor the processes that run on the CPU. You can obtain information about the percentage of CPU that is used by a certain process. CPU usage statistics are sorted in descending order to display the highest consumer at the top. Also included is information about the load on the CPU per process, at 5 seconds, 1 minute, and 5 minutes before the log time. This information is updated automatically every 5 seconds to provide real-time statistics. In ASDM, it is updated every 30 seconds.

To view CPU usage on a per-process basis, choose **Monitoring > Properties > Per-Process CPU Usage**.

You can stop the auto refresh, manually refresh the information, or save it to a file. You can also click the **Configure CPU Usage Colors** button to choose background and foreground colors based on usage percentages, to make it easier to scan for high-usage processes.

## Monitoring Connections

To view current connections in a tabular format, in the ASDM main window, choose **Monitoring > Properties > Connections**. Information for each connection includes the protocol, source and destination address characteristics, idle time since the last packet was sent or received, and the amount of traffic in the connection.





## **PART 8**

### **Logging, SNMP, and Smart Call Home**





## Logging

---

This chapter describes how to log system messages and use them for troubleshooting.

- [About Logging, page 38-1](#)
- [Guidelines for Logging, page 38-5](#)
- [Configure Logging, page 38-6](#)
- [Monitoring the Logs, page 38-24](#)
- [History for Logging, page 38-27](#)

### 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 Cisco 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.

## 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.

## Syslog Message Analysis

The following are some examples of the type of information you can obtain from a review of various syslog messages:

- Connections that are allowed by ASA and ASASM security policies. These messages help you spot holes that remain open in your security policies.
- Connections that are denied by ASA and ASASM security policies. These messages show what types of activity are being directed toward your secured inside network.
- Using the ACE deny rate logging feature shows attacks that are occurring on your ASA or ASA Services Module.
- IDS activity messages can show attacks that have occurred.
- User authentication and command usage provide an audit trail of security policy changes.
- Bandwidth usage messages show each connection that was built and torn down as well as the duration and traffic volume used.
- Protocol usage messages show the protocols and port numbers used for each connection.
- Address translation audit trail messages record NAT or PAT connections being built or torn down, which are useful if you receive a report of malicious activity coming from inside your network to the outside world.

## Syslog Message Format

Syslog messages begin with a percent sign (%) and are structured as follows:

```
%ASA Level Message_number: Message_text
```

Field descriptions are as follows:

|       |                                                                                                                                                    |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| ASA   | The syslog message facility code for messages that are generated by the ASA and ASASM. This value is always ASA.                                   |
| Level | 1 through 7. The level reflects the severity of the condition described by the syslog message—the lower the number, the more severe the condition. |



|                       |                                                                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <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

The following table 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 38-1** 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.

## Syslog Message Filtering

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 Layer 2 peer.

## Sort Messages in the Log Viewers

You can sort messages in all ASDM log viewers (that is, the Real-Time Log Viewer, the Log Buffer Viewer, and the Latest ASDM Syslog Events Viewer). To sort tables by multiple columns, click the header of the first column that you want to sort by, then press and hold down the **Ctrl** key and at the same time, click the headers of the other column(s) that you want to include in the sort order. To sort messages chronologically, select both the date and time columns; otherwise, the messages are sorted only by date (regardless of the time) or only by time (regardless of the date).

When you sort messages in the Real-Time Log Viewer and in the Latest ASDM Syslog Events Viewer, the new messages that come in appear in the sorted order, instead of at the top, as they normally would be. That is, they are mixed in with the rest of the messages.

## 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.

## 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 time stamp and device ID. The syslog server uses the device ID to identify the syslog generator. You can use the **logging device-id** command to generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster.



### Note

To monitor syslog messages from units in a cluster, you must open an ASDM session to each of the units that you want to monitor.

## Guidelines for Logging

This section includes guidelines and limitations that you should review before configuring logging.

### IPv6 Guidelines

Does not support IPv6.

### Additional Guidelines

- The syslog server must run a server program called syslogd. Windows (except for Windows 95 and Windows 98) provides a syslog server as part of its operating system. For Windows 95 and Windows 98, you must obtain a syslogd server from another vendor.
- To view logs generated by the ASA or ASASM, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the ASA and ASASM generate messages but does not save them to a location from which you can view them. You must specify each different logging output destination separately. For example, to designate more than one syslog server as an output destination, specify separate entries in the **Syslog Server** pane for each syslog server.
- Sending syslogs over TCP is not supported on a standby ASA.
- The ASA supports the configuration of 16 syslog servers with the **logging host** command in single context mode. In multiple context mode, the limitation is 4 servers per context.
- The syslog server should be reachable through the ASA and ASASM. You should configure the ASASM 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 does 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 does 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
```

## Configure Logging

This section describes how to configure logging.

- 
- Step 1** Enable logging. See [Enable Logging, page 38-6](#).
- Step 2** Configure an output destination for syslog messages. See [Configure an Output Destination, page 38-7](#).




---

**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.

---

## Enable Logging

To enable logging, perform the following steps:

### Procedure

- 
- Step 1** In ASDM, choose one of the following:

- **Home > Latest ASDM Syslog Messages > Enable Logging**
- **Configuration > Device Management > Logging > Logging Setup**
- **Monitoring > Real-Time Log Viewer > Enable Logging**
- **Monitoring > Log Buffer > Enable Logging**

**Step 2** Check the **Enable logging** check box to turn on logging.

---

## Configure 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.

## Send 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:

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Logging > Logging Setup**.
- Step 2** Check the **Enable logging** check box to turn on logging for the ASA.
- Step 3** Check the **Enable logging on the failover standby unit** check box to turn on logging for the standby ASA, if available.
- Step 4** Check the **Send debug messages as syslogs** check box to redirect all debugging trace output to system logs. The syslog message does not appear on the console if this option is enabled. Therefore, to view debugging messages, you must have logging enabled at the console and have it configured as the destination for the debugging syslog message number and severity level. The syslog message number to use is **711001**. The default severity level for this syslog message is debugging.
- Step 5** Check the **Send syslogs in EMBLEM format** check box to enable EMBLEM format so that it is used for all logging destinations, except syslog servers.
- Step 6** Specify the size of the internal log buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, messages are overwritten unless you save the logs to an FTP server or to internal flash memory. The default buffer size is 4096 bytes. The range is 4096 to 1048576.
- Step 7** To save the buffer content to the FTP server before it is overwritten, check the **Save Buffer To FTP Server** check box. To allow overwriting of the buffer content, uncheck this check box.
- Step 8** Click **Configure FTP Settings** to identify the FTP server and configure the FTP parameters used to save the buffer content.
- Step 9** Check the **Save Buffer To Flash** check box To save the buffer content to internal flash memory before it is overwritten.



**Note** This option is only available in routed or transparent single mode.

- Step 10** Click **Configure Flash Usage** to specify the maximum space to be used in internal flash memory for logging and the minimum free space to be preserved (in KB). Enabling this option creates a directory called “syslog” on the device disk on which messages are stored.



**Note** This option is only available in single routed or transparent mode.

- Step 11** Specify the queue size for system logs that are to be viewed in the ASA or ASASM.

## Configure FTP Settings

To specify the configuration for the FTP server that is used to save the log buffer content, perform the following steps:

### Procedure

- Step 1** Check the **Enable FTP client** check box to enable configuration of the FTP client.
- Step 2** Specify the IP address of the FTP server.
- Step 3** Specify the directory path on the FTP server to store the saved log buffer content.
- Step 4** Specify the username to log in to the FTP server.
- Step 5** Specify the password associated with the username to log in to the FTP server.
- Step 6** Confirm the password, then click **OK**.

## Configure Logging Flash Usage

To specify the limits for saving the log buffer content to internal flash memory, perform the following steps:

### Procedure

- Step 1** Specify the maximum amount of internal flash memory that can be used for logging (in KB).
- Step 2** Specify the amount of internal flash memory that is preserved (in KB). When the internal flash memory approaches that limit, new logs are no longer saved.
- Step 3** Click **OK** to close the **Configure Logging Flash Usage** dialog box.

## Configure Syslog Messaging

To configure syslog messaging, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Setup**.
- Step 2** Choose a system log facility for syslog servers to use as a basis to file messages. The default is LOCAL(4)20, which is what most UNIX systems expect. However, because your network devices share eight available facilities, you might need to change this value for system logs.
- Step 3** Check the **Include timestamp in syslog** check box to add the date and time in each syslog message sent.
- Step 4** Choose the information to be displayed in the **Syslog ID** table. Available options are as follows:
- Choose **Show all syslog IDs** to specify that the **Syslog ID** table should display the entire list of syslog message IDs.
  - Choose **Show disabled syslog IDs** to specify that the **Syslog ID** table should display only those syslog message IDs that have been explicitly disabled.
  - Choose **Show syslog IDs with changed logging** to specify that the **Syslog ID** table should display only those syslog message IDs with severity levels that have changed from their default values.
  - Choose **Show syslog IDs that are disabled or with a changed logging level** to specify that the **Syslog ID** table should display only those syslog message IDs with severity levels that have been modified and the IDs of syslog messages that have been explicitly disabled.
- Step 5** The **Syslog ID Setup Table** displays the list of syslog messages based on the setting in the Syslog ID Setup Table. Choose individual messages or ranges of message IDs that you want to modify. You can either disable the selected message IDs or modify their severity levels. To select more than one message ID in the list, click the first ID in the range and Shift-click the last ID in the range.
- Step 6** Click **Advanced** to configure syslog messages to include a device ID.
- 

## Edit Syslog ID Settings

To change syslog message settings, perform the following steps:



### Note

The **Syslog ID(s)** field is display-only. The values that appear in this area are determined by the entries you chose in the **Syslog ID** table, located in the **Syslog Setup** pane.

---

### Procedure

- 
- Step 1** Check the **Disable Message(s)** check box to disable messages for the syslog message ID(s) displayed in the **Syslog ID(s)** list.
- Step 2** Choose the severity logging level of messages to be sent for the syslog message ID(s) displayed in the **Syslog ID(s)** list. Severity levels are defined as follows:
- Emergency (level 0, system is unusable)



### Note

Using a severity level of zero is not recommended.

---

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

**Step 3** Click **OK** to close the **Edit Syslog ID Settings** dialog box.

---

## Include a Device ID in Non-EMBLEM Formatted Syslog Messages

To include a device ID in non-EMBLEM formatted syslog messages, perform the following steps:

### Procedure

- 
- Step 1** Check the **Enable syslog device ID** check box to specify that a device ID should be included in all non-EMBLEM formatted syslog messages.
- Step 2** To specify which to use as the device ID, choose one of the following options:
- Hostname of the ASA
  - Interface IP address  
Choose the interface name that corresponds to the selected IP address from the drop-down list.  
Check the **In an ASA cluster, always use master's IP address for the selected interface** check box if you are using clustering.
  - String  
Specify an alphanumeric, user-defined string.
  - ASA cluster name
- Step 3** Click **OK** to close the **Advanced Syslog Configuration** dialog box.
- 

## Send 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:

### Procedure

- 
- Step 1** Choose one of the following options to specify which syslog messages should be sent to the internal log buffer:



- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
- **Configuration > Device Management > Logging > Logging Filters**

**Step 2** Choose **Monitoring > Logging > Log Buffer > View**. Then choose **File > Clear Internal Log Buffer** in the **Log Buffer** pane to empty the internal log buffer.

**Step 3** Choose **Configuration > Device Management > Logging > Logging Setup** to change the size of the internal log buffer. The default buffer size is 4 KB.

The ASA and ASASM continue to save new messages to the internal log buffer and save the full log buffer content to internal flash memory. When saving the buffer content to another location, the ASA and ASASM create log files with names that use the following time-stamp format:

*LOG-YYYY-MM-DD-HHMMSS.TXT*

where *YYYY* is the year, *MM* is the month, *DD* is the day of the month, and *HHMMSS* is the time in hours, minutes, and seconds.

**Step 4** To save new messages to another location, choose one of the following options:

- Check the **Flash** check box to send new messages to internal flash memory, then click **Configure Flash Usage**. The **Configure Logging Flash Usage** dialog box appears.
    - a. Specify the maximum amount of flash memory in KB that you want to use for logging.
    - b. Specify the minimum amount of free space in KB that logging will preserve in flash memory.
    - c. Click **OK** to close this dialog box.
  - Check the **FTP Server** check box to send new messages to an FTP server, then click **Configure FTP Settings**. The **Configure FTP Settings** dialog box appears.
    - a. Check the **Enable FTP Client** check box.
    - b. Enter the following information in the fields provided: FTP server IP address, path, username, and password.
    - c. Confirm the password, then click **OK** to close this dialog box.
- 

## Save an Internal Log Buffer to Flash

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

### Procedure

---

**Step 1** Choose **File > Save Internal Log Buffer to Flash**.

The **Enter Log File Name** dialog box appears.

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

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

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

---

## View and Copy Logged Entries with the ASDM Java Console

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

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

### Procedure

- 
- Step 1** Choose **Tools > ASDM Java Console**.
  - Step 2** Enter **m** in the console to show the virtual machine memory statistics.
  - Step 3** Enter **g** in the console to perform garbage collection.
  - Step 4** Open the Windows Task Manager and double-click the **asdm\_launcher.exe** file to monitor memory usage.



**Note** The maximum memory allocation allowed is 256 MB.

---

## Send Syslog Messages to an E-mail Address

To send syslog messages to an e-mail address, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > E-Mail Setup**.
  - Step 2** Specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.
  - Step 3** Click **Add** to enter a new e-mail address recipient of the specified syslog messages.
  - Step 4** Choose the severity level of the syslog messages that are sent to the recipient from the drop-down list. The syslog message severity filter used for the destination e-mail address causes messages of the specified severity level and higher to be sent. The global filter specified in the **Logging Filters** pane is also applied to each e-mail recipient.
  - Step 5** Click **Edit** to modify an existing severity level of the syslog messages that are sent to this recipient.
  - Step 6** Click **OK** to close the **Add E-mail Recipient** dialog box.
- 

## Add or Edit E-Mail Recipients

To add or edit e-mail recipients and severity levels, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > E-mail Setup**.
  - Step 2** Click **Add** or **Edit** to display the **Add/Edit E-Mail Recipient** dialog box.

**Step 3** Enter the destination e-mail address, and choose the syslog severity level from the drop-down list. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)



**Note** Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)



**Note** The severity level used to filter messages for the destination e-mail address is the higher of the severity level specified in the **Add/Edit E-Mail Recipient** dialog box and the global filter set for all e-mail recipients in the **Logging Filters** pane.

**Step 4** Click **OK** to close the **Add/Edit E-Mail Recipient** dialog box.

The added or revised entry appears in the **E-mail Recipients** pane.

**Step 5** Click **Apply** to save your changes to the running configuration.

## Configure the Remote SMTP Server

To configure the remote SMTP server to which e-mail alerts and notifications are sent in response to specific events, perform the following steps:

### Procedure

**Step 1** Choose **Configuration > Device Setup > Logging > SMTP**.

**Step 2** Enter the IP address of the primary SMTP server.

**Step 3** (Optional) Enter the IP address of the standby SMTP server, then click **Apply** to save your changes to the running configuration.

## View Syslog Messages in ASDM

To view the latest syslog messages that have been sent to ASDM, perform the following steps:

### Procedure

**Step 1** Choose **Home > Latest ASDM Syslog Messages**.

The ASA or ASASM sets aside a buffer area for syslog messages waiting to be sent to ASDM and saves messages in the buffer as they occur. The ASDM log buffer is a different buffer than the internal log buffer. When the ASDM log buffer is full, the ASA or ASASM deletes the oldest syslog message to make room in the buffer for new ones. Deleting the oldest syslog message to make room for new ones is the default setting in ASDM.

---

## Apply Message Filters to a Logging Destination

To apply message filters to a logging destination, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Logging Filters**.
- Step 2** Choose the name of the logging destination to which you want to apply a filter. Available logging destinations are as follows:
- ASDM
  - Console port
  - E-Mail
  - Internal buffer
  - SNMP server
  - Syslog server
  - Telnet or SSH session
- Included in this selection are the second column, Syslogs From All Event Classes, and the third column, Syslogs From Specific Event Classes. The second column lists the severity or the event class to use to filter messages for the logging destination, or whether logging is disabled for all event classes. The third column lists the event class to use to filter messages for that logging destination.
- Step 3** Click **Edit** to display the **Edit Logging Filters** dialog box. To apply, edit, or disable filters, see [Apply Logging Filters, page 38-14](#).
- 

## Apply Logging Filters

To apply filters, perform the following steps:

### Procedure

- 
- Step 1** Choose the **Filter on severity** option to filter syslog messages according to their severity level.
- Step 2** Choose the **Use event list** option to filter syslog messages according to an event list.
- Step 3** Choose the **Disable logging from all event classes** option to disable all logging to the selected destination.
- Step 4** Click **New** to add a new event list. To add a new event list, see [Create a Custom Event List, page 38-16](#).
- Step 5** Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

- Step 6** Choose the level of logging messages from the drop-down list. Severity levels include the following:
- Emergency (level 0, system is unusable)



**Note** Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

- Step 7** Click **Add** to add the event class and severity level, then click **OK**.  
The selected logging destination for a filter appears at the top.

## Add or Edit a Message Class and Severity Filter

To add or edit a message class and severity level for filtering messages, perform the following steps:

### Procedure

- Step 1** Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

- Step 2** Choose the level of logging messages from the drop-down list. Severity levels include the following:
- Emergency (level 0, system is unusable)



**Note** Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

- Step 3** Click **OK** when you are done making selections.

## Add or Edit a Syslog Message ID Filter

To add or edit a syslog message ID filter, see [Edit Syslog ID Settings, page 38-9](#).

## Send Syslog Messages to the Console Port

To send syslog messages to the console port, perform the following steps:

### Procedure

- 
- Step 1** Choose one of the following options:
- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
  - **Configuration > Device Management > Logging > Logging Filters**
- Step 2** Select the console in the **Logging Destination** column, then click **Edit**.  
The **Edit Logging Filters** dialog box appears.
- Step 3** Choose either syslogs from all event classes or syslogs from specific event classes to specify which syslog messages should be sent to the console port.
- 

## Send Syslog Messages to a Telnet or SSH Session

To send syslog messages to a Telnet or SSH session, perform the following steps:

### Procedure

- 
- Step 1** Choose one of the following option:
- **Home > Latest ASDM Syslog Messages > Configure ASDM Syslog Filters**
  - **Configuration > Device Management > Logging > Logging Filters**
- Step 2** Select the **Telnet** and **SSH Sessions** in the **Logging Destination** column, then click **Edit**.  
The **Edit Logging Filters** dialog box appears.
- Step 3** Choose either syslogs from all event classes or syslogs from specific event classes to specify which syslog messages should be sent to a Telnet or an SSH session..
- Step 4** Choose **Configuration > Device Management > Logging > Logging Setup** to enable logging for the current session only.
- Step 5** Check the **Enable logging** check box, then click **Apply**.
- 

## Create a Custom Event List

You use the following three criteria to define an event list:

- Event Class
- Severity
- Message ID

To create a custom event list to send to a specific logging destination (for example, an SNMP server), perform the following steps:

#### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Event Lists**.
- Step 2** Click **Add** to display the **Add Event List** dialog box.
- Step 3** Enter the name of the event list. No spaces are allowed.
- Step 4** Click **Add** to display the **Add Class and Severity Filter** dialog box.
- Step 5** Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.
- Step 6** Choose the severity level from the drop-down list. Severity levels include the following:
- Emergency (level 0, system is unusable)



**Note** Using a severity level of zero is not recommended.

---

- Alert (level 1, immediate action is needed)
  - Critical (level 2, critical conditions)
  - Error (level 3, error conditions)
  - Warning (level 4, warning conditions)
  - Notification (level 5, normal but significant conditions)
  - Informational (level 6, informational messages only)
  - Debugging (level 7, debugging messages only)
- Step 7** Click **OK** to close the **Add Event List** dialog box.
- Step 8** Click **Add** to display the **Add Syslog Message ID Filter** dialog box.
- Step 9** Enter a syslog message ID or range of IDs (for example, 101001-199012) to include in the filter.
- Step 10** Click **OK** to close the **Add Event List** dialog box.
- The event of interest appears in the list.
- 

## Generate Syslog Messages in EMBLEM Format to a Syslog Server

To generate syslog messages in EMBLEM format to a syslog server, perform the following steps:

#### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Server**.
- Step 2** Click **Add** to add a new syslog server.
- The **Add Syslog Server** dialog box appears.



---

**Note** You can set up a maximum of four syslog servers per security context (up to a total of 16).

---

- Step 3** Specify the number of messages that are allowed to be queued on the ASA or ASASM when a syslog server is busy. A zero value means an unlimited number of messages may be queued.
- Step 4** Check the **Allow user traffic to pass when TCP syslog server is down** check box to specify whether or not to restrict all traffic if any syslog server is down. If you specify TCP, the ASA or ASASM discovers when the syslog server fails and as a security protection, new connections through the ASA are blocked. If you specify UDP, the ASA or ASASM continues to allow new connections whether or not the syslog server is operational. Valid port values for either protocol are 1025 through 65535. The default UDP port is 514. The default TCP port is 1470.



---

**Note** Sending syslogs over TCP is not supported on a standby ASA.

---

## Add or Edit Syslog Server Settings

To add or edit syslog server settings, perform the following steps:

### Procedure

- 
- Step 1** Choose the interface used to communicate with the syslog server from the drop-down list.
- Step 2** Enter the IP address that is used to communicate with the syslog server.
- Choose the protocol (either TCP or UDP) that is used by the syslog server to communicate with the ASA or ASASM. You can configure the ASA and ASASM to send data to a syslog server using either UDP or TCP, but not both. The default protocol is UDP if you do not specify a protocol.
- Step 3** Enter the port number used by the syslog server to communicate with the ASA or ASASM.
- Step 4** Check the **Log messages in Cisco EMBLEM format (UDP only)** check box to specify whether to log messages in Cisco EMBLEM format (available only if UDP is selected as the protocol).
- Step 5** Check the **Enable secure logging using SSL/TLS (TCP only)** check box to specify that the connection to the syslog server is secure through the use of SSL/TLS over TCP, and that the syslog message content is encrypted.
- Step 6** Click **OK** to complete the configuration.
- 

## Generate Syslog Messages in EMBLEM Format to Other Output Destinations

To generate syslog messages in EMBLEM format to other output destinations, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Logging Setup**.



- Step 2** Check the **Send syslogs in EMBLEM format** check box.
- 

## Change the Amount of Internal Flash Memory Available for Logs

To change the amount of internal flash memory available for logs, perform the following steps:

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Logging > Logging Setup**.
- Step 2** Check the **Enable Logging** check box.
- Step 3** Check the **Save Buffer to Flash** check box in the **Logging to Internal Buffer** area.
- Step 4** Click **Configure Flash Usage**.
- The **Configure Logging Flash Usage** dialog box appears.
- Step 5** Enter the maximum amount of flash memory in KB allowed to be used for logging.
- By default, the ASA can use up to 1 MB of internal flash memory for log data. The minimum amount of internal flash memory that must be free for the ASA and ASASM to save log data is 3 MB. If a log file being saved to internal flash memory would cause the amount of free internal flash memory to fall below the configured minimum limit, the ASA or ASASM deletes the oldest log files to ensure that the minimum amount of memory remains free after saving the new log file. If there are no files to delete or if, after all old files have been deleted, free memory is still below the limit, the ASA or ASASM fails to save the new log file.
- Step 6** Enter the minimum amount of free space in KB to be preserved for logging in flash memory.
- Step 7** Click **OK** to close the **Configure Logging Flash Usage** dialog box.
- 

## Configure the Logging Queue

To configure the logging queue, perform the following steps:

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Logging > Logging Setup**.
- Step 2** Check the **Enable logging** check box.
- Step 3** Enter the number of syslog messages that the ASA and ASASM can hold in its queue before sending them to the configured output destination.

The ASA and ASASM have a fixed number of blocks in memory that can be allocated for buffering syslog messages while they are waiting to be sent to the configured output destination. The number of blocks required depends on the length of the syslog message queue and the number of syslog servers specified. The default queue size is 512 syslog messages. The queue size is limited only by block memory availability. Valid values are from 0 to 8192 messages, depending on the platform. If the logging queue is set to zero, the queue is the maximum configurable size (8192 messages).

- Step 4** Click **Apply** to save your changes to the running configuration.
- 

## Send All Syslog Messages in a Class to a Specified Output Destination

To send all syslog messages in a class to a specified output destination, perform the following steps:

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Logging > Logging Filters**.
- Step 2** To override the configuration in the specified output destination, choose the output destination that you want to change, then click **Edit**.  
The **Edit Logging Filters** dialog box appears.
- Step 3** Revise the settings in either the **Syslogs from All Event Classes** or **Syslogs from Specific Event Classes** area, then click **OK** to close this dialog box.  
For example, if you specify that messages at severity level 7 should go to the internal log buffer and that ha class messages at severity level 3 should go to the internal log buffer, then the latter configuration takes precedence.  
To specify that a class should go to more than one destination, select a different filtering option for each output destination.
- 

## Enable Secure Logging

To enable secure logging, perform the following steps:

### Procedure

---

- Step 1** Choose **Configuration > Device Management > Logging > Syslog Server**.
- Step 2** Select a syslog server for which you want to enable secure logging, then click **Edit**.  
The **Edit Syslog Server** dialog box appears.
- Step 3** Click the **TCP** radio button.



**Note** Secure logging does not support UDP; an error occurs if you try to use this protocol.

---

- Step 4** Check the **Enable secure syslog with SSL/TLS** check box, then click **OK**.
-

## Include the Device ID in Non-EMBLEM Format Syslog Messages

To include the device ID in non-EMBLEM format syslog messages, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration**.
- Step 2** Check the **Enable syslog device ID** check box.
- Step 3** Click the **Hostname**, **Interface IP Address**, or **String** radio button in the **Device ID** area.
- If you chose the **Interface IP Address** option, make sure that the correct interface is selected in the drop-down list.
  - If you chose the **String** option, enter the device ID in the **User-Defined ID** field. The string can include as many as 16 characters.



**Note** If enabled, the device ID does not appear in EMBLEM-formatted syslog messages nor in SNMP traps.

---

- Step 4** Click **OK** to close the **Advanced Syslog Configuration** dialog box.
- 

## Include the Date and Time in Syslog Messages

To include the date and time in syslog messages, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Setup**.
- Step 2** Check the **Include timestamp in syslogs** check box in the **Syslog ID Setup** area.
- Step 3** Click **Apply** to save your changes.
- 

## Disable a Syslog Message

To disable a specified syslog message, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Setup**.
- Step 2** Select the syslog that you want to disable from the table, then click **Edit**.  
The **Edit Syslog ID Settings** dialog box appears.
- Step 3** Check the **Disable messages** check box, then click **OK**.
-

## Change the Severity Level of a Syslog Message

To change the severity level of a syslog message, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Setup**.
  - Step 2** Select the syslog whose severity level you want to change from the table, then click **Edit**.  
The **Edit Syslog ID Settings** dialog box appears.
  - Step 3** Choose the desired severity level from the **Logging Level** drop-down list, then click **OK**.
- 

## Block Syslog Messages on a Standby Unit

To block specific syslog messages from being generated on a standby unit, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Syslog Settings**.
  - Step 2** Choose a syslog ID in the table, then click **Edit**.  
The **Edit Syslog ID Settings** dialog box appears.
  - Step 3** Check the **Disable messages on standby unit** check box to block syslog messages from being generated on a standby unit.
  - Step 4** Click **OK** to close this dialog box.
- 

## Limit the Rate of Syslog Message Generation

To limit the rate of syslog message generation, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Logging > Rate Limit**.
  - Step 2** Choose the logging level (message severity level) to which you want to assign rate limits. Severity levels are defined as follows:

| Description | Severity Level               |
|-------------|------------------------------|
| Emergency   | 0—System is unusable         |
| Alert       | 1—Immediate action is needed |
| Critical    | 2—Critical conditions        |
| Error       | 3—Error conditions           |
| Warning     | 4—Warning conditions         |

| Description   | Severity Level                      |
|---------------|-------------------------------------|
| Notification  | 5—Normal but significant conditions |
| Informational | 6—Informational messages only       |
| Debugging     | 7—Debugging messages only           |

- Step 3** The No of Messages field displays the number of messages sent. The Interval (Seconds) field displays the interval, in seconds, that is used to limit how many messages at this logging level can be sent. Choose a logging level from the table and click **Edit** to display the **Edit Rate Limit for Syslog Logging Level** dialog box.
- Step 4** To continue, see [Assign or Change Rate Limits for Individual Syslog Messages, page 38-23](#).

## Assign or Change Rate Limits for Individual Syslog Messages

To assign or change rate limits to individual syslog messages, perform the following steps:

### Procedure

- Step 1** To assign the rate limit of a specific syslog message, click **Add** to display the **Add Rate Limit for Syslog Message** dialog box.
- Step 2** To continue, see [Add or Edit the Rate Limit for a Syslog Message, page 38-23](#).
- Step 3** To change the rate limit of a specific syslog message, click **Edit** to display the **Edit Rate Limit for Syslog Message** dialog box.
- Step 4** To continue, see [Edit the Rate Limit for a Syslog Severity Level, page 38-24](#).

## Add or Edit the Rate Limit for a Syslog Message

To add or change the rate limit for a specific syslog message, perform the following steps:

### Procedure

- Step 1** To add a rate limit to a specific syslog message, click **Add** to display the **Add Rate Limit for Syslog Message** dialog box. To change a rate limit for a syslog message, click **Edit** to display the **Edit Rate Limit for Syslog Message** dialog box.
- Step 2** Enter the message ID of the syslog message that you want to limit.
- Step 3** Enter the maximum number of messages that can be sent in the specified time interval.
- Step 4** Enter the amount of time, in seconds, that is used to limit the rate of the specified message, then click **OK**.

**Note**

To allow an unlimited number of messages, leave both the **Number of Messages** and **Time Interval** fields blank.

## Edit the Rate Limit for a Syslog Severity Level

To change the rate limit of a specified syslog severity level, perform the following steps:

**Procedure**

- 
- Step 1** Enter the maximum number of messages at this severity level that can be sent.
- Step 2** Enter the amount of time, in seconds, that is used to limit the rate of messages at this severity level, and click **OK**.

The selected message severity level appears.

**Note**

To allow an unlimited number of messages, leave both the **Number of Messages** and **Time Interval** fields blank.

## Monitoring the Logs

See the following screens for monitoring logging status.

- **Monitoring > Logging > Log Buffer > View**  
This pane allows you to view the log buffer.
- **Monitoring > Logging > Real-Time Log Viewer > View**  
This pane allows you to view the real-time log.
- **Tools > Command Line Interface**  
This pane allows you to issue various non-interactive commands and view results.

## Filter Syslog Messages Through the Log Viewers

You can filter syslog messages based on one or multiple values that correspond to any column of the Real-Time Log Viewer and the Log Buffer Viewer.

To filter syslog messages through one of the log viewers, perform the following steps:

**Procedure**

- 
- Step 1** Choose one of the following options:
- **Monitoring > Logging > Real-Time Log Viewer > View**

- **Monitoring > Logging > Log Buffer > View**

**Step 2** In either the **Real-Time Log Viewer** or the **Log Buffer Viewer** dialog box, click **Build Filter** on the toolbar.

**Step 3** In the **Build Filter** dialog box, specify the filtering criteria to apply to syslog messages:

- Choose one of the following three options in the **Date and Time** area: real-time, a specific time, or a time range. If you chose a specific time, indicate the time by entering the number and choosing hours or minutes from the drop-down list. If you chose a time range, click the drop-down arrow in the **Start Time** field to display a calendar. Choose a start date and a start time from the drop-down list, then click **OK**. Click the drop-down arrow in the **End Time** field to display a calendar. Choose an end date and an end time from the drop-down list, then click **OK**.
- Enter a valid severity level in the **Severity** field. Alternatively, click the **Edit** icon on the right of the **Severity** field. Click the severity levels in the list on which you want to filter. To include severity levels 1-7, click **All**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Severity** field for additional information about the correct input format to use.
- Enter a valid syslog ID in the **Syslog ID** field. Alternatively, click the **Edit** icon on the right of the **Syslog ID** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Syslog ID** field for additional information about the correct input format to use.
- Enter a valid source IP address in the **Source IP Address** field, or click the **Edit** icon on the right of the **Source IP Address** field. Choose a single IP address or a specified range of IP addresses, then click **Add**. Check the **Do not include (exclude) this address or range** check box to exclude a specific IP address or range of IP addresses. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Source IP Address** field for additional information about the correct input format to use.
- Enter a valid source port in the **Source Port** field, or click the **Edit** icon on the right of the **Source Port** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Source Port** field for additional information about the correct input format to use.
- Enter a valid destination IP address in the **Destination IP Address** field, or click the **Edit** icon on the right of the **Destination IP Address** field. Choose a single IP address or a specified range of IP addresses, then click **Add**. Check the **Do not include (exclude) this address or range** check box to exclude a specific IP address or range of IP addresses. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Destination IP Address** field for additional information about the correct input format to use.
- Enter a valid destination port in the **Destination Port** field, or click the **Edit** icon on the right of the **Destination Port** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Destination Port** field for additional information about the correct input format to use.
- Enter filtering text for the **Description** field. The text may be any string of one or more characters, including a regular expression. However, semicolons are not valid characters, and this setting is case-sensitive. Multiple entries must be separated by commas.
- Click **OK** to add the filter settings you have just specified to the **Filter By** drop-down list in the log viewers. The filter strings follow a specific format. The prefix **FILTER:** designates all custom filters that appear in the **Filter By** drop-down list. You may still type random text into this field.

The following table shows examples of the format used.

| Build Filter Example                                                  | Filter String Format                          |
|-----------------------------------------------------------------------|-----------------------------------------------|
| Source IP = 192.168.1.1 or 0.0.0.0<br>Source Port = 67                | FILTER: srcIP=192.168.1.1,0.0.0.0;srcPort=67; |
| Severity = Informational<br>Destination IP = 1.1.1.1 through 1.1.1.10 | FILTER: sev=6;dstIP=1.1.1.1-1.1.1.10;         |
| Syslog ID not in the range 725001 through 725003                      | FILTER: sysID=!725001-725003;                 |
| Source IP = 1.1.1.1<br>Description = Built outbound                   | FILTER: srcIP=1.1.1.1;descr=Built outbound    |

- Step 4** Choose one of the settings in the **Filter By** drop-down list to filter syslog messages, then click **Filter** on the toolbar. This setting also applies to all future syslog messages. Click **Show All** on the toolbar to clear all filters.



**Note** You cannot save filters that you have specified with the **Build Filter** dialog box. These filters are valid only for the ASDM session during which they were created.

## Edit Filtering Settings

To edit filtering settings that you created using the **Build Filter** dialog box, perform the following steps:

### Procedure

- Step 1** Choose one of the following options:
- Revise a filter directly by entering the changes in the **Filter By** drop-down list.
  - Choose a filter in the **Filter By** drop-down list, then click **Build Filter** to display the **Build Filter** dialog box. Click **Clear Filter** to remove the current filter settings and enter new ones. Otherwise, change the settings that appear, and click **OK**.



**Note** These filter settings apply only to those defined in the **Build Filter** dialog box.

- Click **Show All** on the toolbar to stop filtering and show all syslog messages.



## Issue Certain Commands Using the Log Viewers

You can issue the following commands using either of the log viewers: **ping**, **traceroute**, **whois**, and **dns lookup**.

To run any of these commands, perform the following steps:

### Procedure

- 
- Step 1** Choose one of the following options:
- **Monitoring > Logging > Real-Time Log Viewer > View**
  - **Monitoring Logging > Log Buffer > View**
- Step 2** Click **Tools** from the **Real-Time Log Viewer** or **Log Buffer** pane, then choose the command that you want to execute. Alternatively, you can right-click a specific syslog message that is listed to display a context menu, then choose the command that you want to execute.
- The **Entering command** dialog box appears, with the command that you selected automatically showing in the drop-down list.
- Step 3** Enter either the source or destination IP address of the selected syslog message in the **Address** field, then click **Go**.
- The command output appears in the area provided.
- Step 4** Click **Clear** to remove the output, and choose another command to execute from the drop-down list. Repeat Step 3, if necessary. Click **Close** when you are done.
- 

## History for Logging

**Table 38-2** History for Logging

| Feature Name | Platform Releases | Description                                                                                                                                                                                                                                |
|--------------|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logging      | 7.0(1)            | Provides ASA network logging information through various output destinations, and includes the option to view and save log files.<br><br>We introduced the following screen: Configuration > Device Management > Logging > Logging Setup.  |
| Rate limit   | 7.0(4)            | Limits the rate at which syslog messages are generated.<br><br>We modified the following screen: Configuration > Device Management > Logging > Rate Limit.                                                                                 |
| Logging list | 7.2(1)            | Creates a logging list to use in other commands to specify messages by various criteria (logging level, event class, and message IDs).<br><br>We modified the following screen: Configuration > Device Management > Logging > Event Lists. |

**Table 38-2** History for Logging (continued)

| Feature Name                             | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.<br><br>We modified the following screen: Configuration > Device Management > Logging > Syslog Server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Logging class                            | 8.0(4), 8.1(1)    | Added support for the ipaa event class of logging messages.<br><br>We modified the following screen: Configuration > Device Management > Logging > Logging Filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Logging class and saved logging buffers  | 8.2(1)            | Added support for the dap event class of logging messages.<br><br>Added support to clear the saved logging buffers (ASDM, internal, FTP, and flash).<br><br>We modified the following screen: Configuration > Device Management > Logging > Logging Setup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Password encryption                      | 8.3(1)            | Added support for password encryption.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Log viewers                              | 8.3(1)            | The source and destination IP addresses were added to the log viewers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 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.<br><br>This feature was added for compliance with Common Criteria EAL4+. Unless required, we recommended allowing connections when syslog messages cannot be sent or received. To allow connections, continue to check the <b>Allow user traffic to pass when TCP syslog server is down</b> check box on the Configuration > Device Management > Logging > Syslog Servers pane.<br><br>We introduced the following syslog messages: 414005, 414006, 414007, and 414008.<br><br>We did not modify any ASDM screens. |
| Syslog message filtering and sorting     | 8.4(1)            | Support has been added for the following: <ul style="list-style-type: none"> <li>• Syslog message filtering based on multiple text strings that correspond to various columns</li> <li>• Creation of custom filters</li> <li>• Column sorting of messages. For detailed information, see the ASDM configuration guide.</li> </ul> We modified the following screens:<br>Monitoring > Logging > Real-Time Log Viewer > View.<br>Monitoring > Logging > Log Buffer Viewer > View.<br><br>This feature interoperates with all ASA versions.                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 38-2**      *History for Logging (continued)*

| Feature Name                       | Platform Releases | Description                                                                                                                                                                                                                     |
|------------------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Clustering                         | 9.0(1)            | Added support for syslog message generation in a clustering environment on the ASA 5580 and 5585-X.<br><br>We modified the following screen: Configuration > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration. |
| Blocking syslogs on a standby unit | 9.4(1)            | We added support for blocking the generation of specific syslog messages on the standby unit in a failover configuration.<br><br>We modified the following screen: Configuration > Device Management > Logging > Syslog Setup.  |





# SNMP

---

This chapter describes how to configure Simple Network Management Protocol (SNMP) to monitor the Cisco ASA.

- [About SNMP, page 39-1](#)
- [Guidelines for SNMP, page 39-4](#)
- [Configure SNMP, page 39-5](#)
- [Monitoring SNMP, page 39-10](#)
- [History for SNMP, page 39-10](#)

## 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. The ASA, ASAv, and ASASM provide support for network monitoring using SNMP Versions 1, 2c, and 3, and support 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, ASAv, 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, ASAv, 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 Management Information Bases (MIBs) on the ASA. MIBs are a collection of definitions, and the ASA, ASAv, 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, ASAv, 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, ASAv, or ASASM SNMP agent also replies when a management station asks for information.

## SNMP Terminology

The following table lists the terms that are commonly used when working with SNMP.

**Table 39-1** *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"><li>• Responds to requests for information and actions from the network management station.</li><li>• Controls access to its Management Information Base, the collection of objects that the SNMP manager can view or change.</li><li>• Does not allow SET operations.</li></ul> |
| 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, ASAv, and ASASM.                                                                                                                                                                                                                                                                                       |
| Object identifier (OID)             | The system that identifies a device to its NMS and indicates to users the source of information monitored and displayed.                                                                                                                                                                                                                                                                          |
| Trap                                | Predefined events that generate a message from the SNMP agent to the NMS. Events include alarm conditions such as linkup, linkdown, coldstart, warmstart, authentication, or syslog messages.                                                                                                                                                                                                     |

## SNMP Version 3 Overview

SNMP Version 3 provides security enhancements that are not available in SNMP Version 1 or Version 2c. SNMP Versions 1 and 2c transmit data between the SNMP server and SNMP agent in clear text. SNMP Version 3 adds authentication and privacy options to secure protocol operations. In addition, this version controls access to the SNMP agent and MIB objects through the User-based Security Model (USM) and View-based Access Control Model (VACM). The ASA and ASASM also support the creation of SNMP groups and users, as well as hosts, which is required to enable transport authentication and encryption for secure SNMP communications.

## Security Models

For configuration purposes, the authentication and privacy options are grouped together into security models. Security models apply to users and groups, which are divided into the following three types:

- NoAuthPriv—No Authentication and No Privacy, which means that no security is applied to messages.
- AuthNoPriv—Authentication but No Privacy, which means that messages are authenticated.
- AuthPriv—Authentication and Privacy, which means that messages are authenticated and encrypted.

## SNMP Groups

An SNMP group is an access control policy to which users can be added. Each SNMP group is configured with a security model, and is associated with an SNMP view. A user within an SNMP group must match the security model of the SNMP group. These parameters specify what type of authentication and privacy a user within an SNMP group uses. Each SNMP group name and security model pair must be unique.

## SNMP Users

SNMP users have a specified username, a group to which the user belongs, authentication password, encryption password, and authentication and encryption algorithms to use. The authentication algorithm options are MD5 and SHA. The encryption algorithm options are DES, 3DES, and AES (which is available in 128, 192, and 256 versions). When you create a user, you must associate it with an SNMP group. The user then inherits the security model of the group.

## SNMP Hosts

An SNMP host is an IP address to which SNMP notifications and traps are sent. To configure SNMP Version 3 hosts, along with the target IP address, you must configure a username, because traps are only sent to a configured user. SNMP target IP addresses and target parameter names must be unique on the ASA and ASA Services Module. Each SNMP host can have only one username associated with it. To receive SNMP traps, configure the SNMP NMS, and make sure that you configure the user credentials on the NMS to match the credentials for the ASA and ASASM.

## Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software

The SNMP Version 3 implementation in the ASA and ASASM differs from the SNMP Version 3 implementation in the Cisco IOS software in the following ways:

- The local-engine and remote-engine IDs are not configurable. The local engine ID is generated when the ASA or ASASM starts or when a context is created.
- No support exists for view-based access control, which results in unrestricted MIB browsing.
- Support is restricted to the following MIBs: USM, VACM, FRAMEWORK, and TARGET.
- You must create users and groups with the correct security model.
- You must remove users, groups, and hosts in the correct sequence.
- Use of the **snmp-server host** command creates an ASA, ASAv, or ASASM rule to allow incoming SNMP traffic.

## SNMP Syslog Messaging

SNMP generates detailed syslog messages that are numbered 212nnn. 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).

## 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](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](http://www.cisco.com/en/US/docs/security/asa/asa83/snmp/snmpv3_tools.html)

## Guidelines for SNMP

This section includes the guidelines and limitations that you should review before configuring SNMP.

### Failover Guidelines

The SNMP client in each ASA, ASAv, 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

- You must have Cisco Works for Windows or another SNMP MIB-II compliant browser to receive SNMP traps or browse a MIB.
- 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 on the ASA 5585 SSP-40 (NPE).
- If the NMS cannot successfully request objects or is not correctly handling incoming traps from the ASA, performing a packet capture is the most useful method for determining the problem. Choose **Wizards > Packet Capture Wizard**, and follow the on-screen instructions.
- You can add up to 4000 hosts. However, only 128 of this number can be for traps.
- The total number of 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.
- You can specify overlapping network objects in different **host-group** commands. The values that you specify for the last host group take effect for the common set of hosts in the different network objects.
- If you delete a host group or hosts that overlap with other host groups, the hosts are set up again using the values that have been specified in the configured host groups.
- The values that the hosts acquire depend on the specified sequence that you use to run the commands.
- The limit on the message size that SNMP sends is 1472 bytes.
- Members of a cluster do not synchronize their SNMPv3 engine IDs. Because of this, each unit in the cluster should have a unique SNMPv3 user configuration.
- With Version 9.4(1), the ASA supports an unlimited number of SNMP server trap hosts per context. The **show snmp-server host** command output displays only the active hosts that are polling the ASA, as well as the statically configured hosts.
- 

## Configure SNMP

This section describes how to configure SNMP.

- 
- |               |                                                                                                                                                                                                                               |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Enable the SNMP Agent and SNMP server. See <a href="#">Enable the SNMP Agent and SNMP Server, page 39-6</a> .                                                                                                                 |
| <b>Step 2</b> | Configure an SNMP management station to receive requests from the ASA. See <a href="#">Configure an SNMP Management Station, page 39-6</a> .                                                                                  |
| <b>Step 3</b> | Configure SNMP traps. See <a href="#">Configure SNMP Traps, page 39-6</a> .                                                                                                                                                   |
| <b>Step 4</b> | Configure SNMP Version 1 and 2c parameters or SNMP Version 3 parameters. See <a href="#">Configure Parameters for SNMP Version 1 or 2c, page 39-7</a> or <a href="#">Configure Parameters for SNMP Version 3, page 39-8</a> . |
-

## Enable the SNMP Agent and SNMP Server

To enable the SNMP agent and SNMP server, perform the following steps:

### Procedure

## Configure an SNMP Management Station

To configure an SNMP management station, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Management Access > SNMP**. By default, the SNMP server is enabled.
  - Step 2** Click **Add** in the **SNMP Management Stations** pane.  
The **Add SNMP Host Access Entry** dialog box appears.
  - Step 3** Choose the interface on which the SNMP host resides.
  - Step 4** Enter the SNMP host IP address.
  - Step 5** Enter the SNMP host UDP port, or keep the default, port 162.
  - Step 6** Add the SNMP host community string. If no community string is specified for a management station, the value set in the **Community String** (default) field on the **SNMP Management Stations** pane is used.
  - Step 7** Choose the SNMP version used by the SNMP host.
  - Step 8** If you have selected SNMP Version 3 in the previous step, choose the name of a configured user.
  - Step 9** To specify the method for communicating with this NMS, check either the **Poll** or **Trap** check box.
  - Step 10** Click **OK**.  
The **Add SNMP Host Access Entry** dialog box closes.
  - Step 11** Click **Apply**.  
The NMS is configured and changes are saved to the running configuration. For more information about SNMP Version 3 NMS tools, see the following URL:  
[http://www.cisco.com/en/US/docs/security/asa/asa82/snmp/snmpv3\\_tools.html](http://www.cisco.com/en/US/docs/security/asa/asa82/snmp/snmpv3_tools.html)
- 

## Configure SNMP Traps

To designate which traps that the SNMP agent generates and how they are collected and sent to NMSs, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Management Access > SNMP**.
  - Step 2** Click **Configure Traps**.  
The **SNMP Trap Configuration** dialog box appears.

**Step 3** Check the **SNMP Server Traps Configuration** check box.

The traps are divided into the following categories: standard, IKEv2, entity MIB, IPsec, remote access, resource, NAT, syslog, CPU utilization, CPU utilization and monitoring interval, and SNMP interface threshold and interval. Check the applicable check boxes for the SNMP events to notify through SNMP traps. The default configuration has all SNMP standard traps enabled. If you do not specify a trap type, the default is the syslog trap. The default SNMP traps continue to be enabled with the syslog trap. All other traps are disabled by default. To disable a trap, uncheck the applicable check box. To configure the syslog trap severity level, choose **Configuration > Device Management > Logging > Logging Filters**.

**Step 4** Click **OK** to close the **SNMP Trap Configuration** dialog box.

**Step 5** Click **Apply**.

The SNMP traps are configured and the changes are saved to the running configuration.

---

## Configure Parameters for SNMP Version 1 or 2c

To configure parameters for SNMP Version 1 or 2c, perform the following steps:

### Procedure

---

**Step 1** Choose **Configuration > Device Management > Management Access > SNMP**.

**Step 2** Enter a default community string in the **Community String** (default) field if you are using SNMP Version 1 or 2c. Enter the password used by the SNMP NMSs when they send requests to the ASA. The SNMP community string is a shared secret among the SNMP NMSs and the network nodes being managed. The ASA uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive value up to 32 alphanumeric characters long. Spaces are not permitted. The default is public. SNMP Version 2c allows separate community strings to be set for each NMS. If no community string is configured for any NMS, the value set here is used by default.

**Step 3** Enter the name of the ASA system administrator. The text is case-sensitive and can be up to 127 alphabetic characters. Spaces are accepted, but multiple spaces are shortened to a single space.

**Step 4** Enter the location of the ASA being managed by SNMP. The text is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.

**Step 5** Enter the number of the ASA port that listens for SNMP requests from NMSs; or keep the default, port number 161.

**Step 6** Click **Add** in the **SNMP Host Access List** pane.

The **Add SNMP Host Access Entry** dialog box appears.

**Step 7** Choose the interface name from which traps are sent from the drop-down list.

**Step 8** Enter the IP address of the NMS or SNMP manager that can connect to the ASA.

**Step 9** Enter the UDP port number. The default is 162.

**Step 10** Choose the SNMP version that you are using from the drop-down list. If you choose Version 1 or Version 2c, you must enter the community string. If you choose Version 3, you must choose the username from the drop-down list.

**Step 11** Check the **Poll** check box in the **Server Poll/Trap Specification** area to limit the NMS to sending requests (polling) only. Check the **Trap** check box to limit the NMS to receiving traps only. You may check both check boxes to perform both functions of the SNMP host.

**Step 12** Click **OK** to close the **Add SNMP Host Access Entry** dialog box.

The new host appears in the **SNMP Host Access List** pane.

**Step 13** Click **Apply**.

SNMP parameters for Versions 1, 2c, or 3 are configured and the changes are saved to the running configuration.

## Configure Parameters for SNMP Version 3

To configure parameters for SNMP Version 3, perform the following steps:

### Procedure

**Step 1** Choose **Configuration > Device Management > Management Access > SNMP**.

**Step 2** Click **Add > SNMP User** on the **SNMPv3 User/Group** tab in the **SNMPv3 Users** pane to add a configured user or a new user to a group. When you remove the last user in a group, ASDM deletes the group.



**Note** After a user has been created, you cannot change the group to which the user belongs.

The **Add SNMP User Entry** dialog box appears.

**Step 3** Choose the group to which the SNMP user belongs. The available groups are as follows:

- **Auth&Encryption**, in which users have authentication and encryption configured
- **Authentication\_Only**, in which users have only authentication configured
- **No\_Authentication**, in which users have neither authentication nor encryption configured



**Note** You cannot change the group names.

**Step 4** Click the **USM Model** tab to use the user security model (USM) groups.

**Step 5** Click **Add**.

The **Add SNMP USM Entry** dialog box appears.

**Step 6** Enter the group name.

**Step 7** Choose the security level from the drop-down list. This setting allows you to assign a configured USM group as a security level to SNMPv3 users.

**Step 8** Enter the name of a configured user or a new user. The username must be unique for the SNMP server group selected.

**Step 9** Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.

**Step 10** Indicate the type of authentication you want to use by clicking one of the two radio buttons: **MD5** or **SHA**.

**Step 11** Enter the password to use for authentication.

- Step 12** Indicate the type of encryption you want to use by clicking one of these three radio buttons: **DES**, **3DES**, or **AES**.
- Step 13** If you chose AES encryption, then choose the level of AES encryption to use: **128**, **192**, or **256**.
- Step 14** Enter the password to use for encryption. The maximum number of alphanumeric characters allowed for this password is 64.
- Step 15** Click **OK** to create a group (if this is the first user in that group), display this group in the **Group Name** drop-down list, and create a user for that group.  
The **Add SNMP User Entry** dialog box closes.
- Step 16** Click **Apply**.  
SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.
- 

## Configure a Group of Users

To configure an SNMP user list with a group of specified users in it, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Management Access > SNMP**.
- Step 2** Click **Add > SNMP User Group** on the **SNMPv3 User/Group** tab in the **SNMPv3 Users** pane to add a configured user group or a new user group. When you remove the last user in a group, ASDM deletes the group.  
The **Add SNMP User Group** dialog box appears.
- Step 3** Enter the user group name.
- Step 4** Click the **Existing User/User Group** radio button to select an existing user or user group.
- Step 5** Click the **Create new user** radio button to create a new user.
- Step 6** Choose the group to which the SNMP user belongs. The available groups are as follows:
- **Auth&Encryption**, in which users have authentication and encryption configured
  - **Authentication\_Only**, in which users have only authentication configured
  - **No\_Authentication**, in which users have neither authentication nor encryption configured
- Step 7** Enter the name of a configured user or a new user. The username must be unique for the SNMP server group selected.
- Step 8** Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.
- Step 9** Indicate the type of authentication you want to use by clicking one of the two radio buttons: **MD5** or **SHA**.
- Step 10** Enter the password to use for authentication.
- Step 11** Confirm the password to use for authentication.
- Step 12** Indicate the type of encryption you want to use by clicking one of these three radio buttons: **DES**, **3DES**, or **AES**.

- Step 13** Enter the password to use for encryption. The maximum number of alphanumeric characters allowed for this password is 64.
- Step 14** Confirm the password to use for encryption.
- Step 15** Click **Add** to add the new user to the specified user group in the **Members in Group** pane. Click **Remove** to delete an existing user from the **Members in Group** pane.
- Step 16** Click **OK** to create a new user for the specified user group.  
The **Add SNMP User Group** dialog box closes.
- Step 17** Click **Apply**.  
SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.

## Monitoring SNMP

See the following screens for monitoring SNMP:

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

## History for SNMP

**Table 39-2** History for SNMP

| Feature Name           | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SNMP Versions 1 and 2c | 7.0(1)            | Provides ASA, ASAv, and ASASM network monitoring and event information by transmitting data between the SNMP server and SNMP agent through the clear text community string.<br><br>We modified the following screen: Configuration > Device Management > Management Access > SNMP.                                                                                                                                                                                         |
| SNMP Version 3         | 8.2(1)            | Provides 3DES or AES encryption and support for SNMP Version 3, the most secure form of the supported security models. This version allows you to configure users, groups, and hosts, as well as authentication characteristics by using the USM. In addition, this version allows access control to the agent and MIB objects and includes additional MIB support.<br><br>We modified the following screen: Configuration > Device Management > Management Access > SNMP. |
| Password encryption    | 8.3(1)            | Supports password encryption.                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 39-2 History for SNMP (continued)

| Feature Name                | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SNMP traps and MIBs         | 8.4(1)            | <p>Supports the following additional keywords: <b>connection-limit-reached</b>, <b>cpu threshold rising</b>, <b>entity cpu-temperature</b>, <b>entity fan-failure</b>, <b>entity power-supply</b>, <b>ikev2 stop   start</b>, <b>interface-threshold</b>, <b>memory-threshold</b>, <b>nat packet-discard</b>, <b>warmstart</b>.</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 modified the following screen: Configuration &gt; Device Management &gt; Management Access &gt; SNMP.</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"> <li>CISCO-ENTITY-SENSOR-EXT-MIB (Only objects under the entPhySensorTable group are supported).</li> <li>ENTITY-SENSOR-MIB (Only objects in the entPhySensorTable group are supported).</li> <li>DISMAN-EXPRESSION-MIB (Only objects in the expExpressionTable, expObjectTable, and expValueTable groups are supported).</li> </ul> <p>Unsupported traps in 8.5(1):</p> <ul style="list-style-type: none"> <li>ceSensorExtThresholdNotification (CISCO-ENTITY-SENSOR-EXT-MIB). This trap is only used for power supply failure, fan failure, and high CPU temperature events.</li> <li>InterfacesBandwidthUtilization.</li> </ul>                                                                                                                                                                                                    |
| 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: <b>entity power-supply-presence</b>, <b>entity power-supply-failure</b>, <b>entity chassis-temperature</b>, <b>entity chassis-fan-failure</b>, <b>entity power-supply-temperature</b>.</p> <p>We modified the following command: <b>snmp-server enable traps</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 39-2 History for SNMP (continued)

| Feature Name                            | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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"> <li>• ALTIGA-GLOBAL-REG.my</li> <li>• ALTIGA-LBSSF-STATS-MIB.my</li> <li>• ALTIGA-MIB.my</li> <li>• ALTIGA-SSL-STATS-MIB.my</li> <li>• CISCO-IPSEC-FLOW-MONITOR-MIB.my</li> <li>• CISCO-REMOTE-ACCESS-MONITOR-MIB.my</li> </ul> |
| 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 <b>show xlate count</b> command.                                                                                                                                                                                                                                                    |
| SNMP hosts, host groups, and user lists | 9.1(5)            | <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 modified the following screen: Configuration &gt; Device Management &gt; Management Access &gt; SNMP.</p>                                                                                           |
| SNMP message size                       | 9.2(1)            | The limit on the message size that SNMP sends has been increased to 1472 bytes.                                                                                                                                                                                                                                                                                                                                                                                            |
| SNMP OIDs and MIBs                      | 9.2(1)            | <p>The ASA now supports the cpmCPUTotal5minRev OID.</p> <p>The ASAv has been added as a new product to the SNMP sysObjectID OID and entPhysicalVendorType OID.</p> <p>The CISCO-PRODUCTS-MIB and CISCO-ENTITY-VENDORTYPE-OID-MIB have been updated to support the new ASAv platform.</p> <p>A new SNMP MIB for monitoring VPN shared license usage has been added.</p>                                                                                                     |
| SNMP OIDs and MIBs                      | 9.3(1)            | CISCO-REMOTE-ACCESS-MONITOR-MIB (OID 1.3.6.1.4.1.9.9.392) support has been added for the ASASM.                                                                                                                                                                                                                                                                                                                                                                            |



Table 39-2 History for SNMP (continued)

| Feature Name                                 | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SNMP MIBs and traps                          | 9.3(2)            | <p>The CISCO-PRODUCTS-MIB and CISCO-ENTITY-VENDORTYPE-OID-MIB have been updated to support the ASA 5506-X.</p> <p>The ASA 5506-X has been added as new products to the SNMP sysObjectID OID and entPhysicalVendorType OID tables.</p> <p>The ASA now supports the CISCO-CONFIG-MAN-MIB, which enables you to do the following:</p> <ul style="list-style-type: none"> <li>• Know which commands have been entered for a specific configuration.</li> <li>• Notify the NMS when a change has occurred in the running configuration.</li> <li>• Track the time stamps associated with the last time that the running configuration was changed or saved.</li> <li>• Track other changes to commands, such as terminal details and command sources.</li> </ul> <p>We modified the following screen: Configuration &gt; Device Management &gt; Management Access &gt; SNMP &gt; Configure Traps &gt; SNMP Trap Configuration.</p> |
| SNMP MIBs and traps                          | 9.4(1)            | <p>The ASA 5506W-X, ASA 5506H-X, ASA 5508-X, and ASA 5516-X have been added as a new product to the SNMP sysObjectID OID and entPhysicalVendorType OID tables.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Unlimited SNMP server trap hosts per context | 9.4(1)            | <p>The ASA supports unlimited SNMP server trap hosts per context. The <b>show snmp-server host</b> command output displays only the active hosts that are polling the ASA, as well as the statically configured hosts.</p> <p>We did not modify any ASDM screens.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |





## Anonymous Reporting and Smart Call Home

This chapter describes how to configure the Anonymous Reporting and Smart Call Home services.

- [About Anonymous Reporting, page 40-1](#)
- [About Smart Call Home, page 40-2](#)
- [Guidelines for Anonymous Reporting and Smart Call Home, page 40-3](#)
- [Configure Anonymous Reporting and Smart Call Home, page 40-4](#)
- [Monitoring Anonymous Reporting and Smart Call Home, page 40-7](#)
- [History for Anonymous Reporting and Smart Call Home, page 40-7](#)

### About Anonymous Reporting

You can help to improve the Cisco 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 the ASA to reach the Cisco Smart Call Home server and send messages to Cisco. Because it is possible that the 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 the 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.

See the syslog messages guide for information about syslog messages.

## 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 on the seriousness of these problems, Cisco responds to your system configuration issues, product end-of-life announcements, security advisory issues, and so on by providing the following services:

- 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.
- Using staff resources more efficiently by reducing troubleshooting time.
- Generating 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 access to required information that enables you to do the following:

- Review all Smart Call Home messages, diagnostics, and recommendations in one place.
- Check service request status.
- View the most up-to-date inventory and configuration information for all Smart Call Home-enabled devices.

# Guidelines for Anonymous Reporting and Smart Call Home

This section includes the guidelines and limitation that you should review before configuring Anonymous reporting and Smart Call Home.

## Anonymous Reporting Guidelines

- DNS must be configured.
- If an Anonymous Reporting message cannot be sent on the first try, the ASA retries two more times before dropping the message.
- Anonymous Reporting may coexist with other Smart Call Home configurations without changing the existing configuration. For example, if Smart Call Home is disabled before enabling Anonymous Reporting, it remains disabled, even after Anonymous Reporting has been enabled.
- 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.

## Smart Call Home Guidelines

- 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

## Related Topics

- [DNS Requirement, page 40-2](#)
- [Configure the DNS Server, page 18-10](#)

# Configure Anonymous Reporting and Smart Call Home

While Anonymous Reporting is part of the Smart Call Home service and allows Cisco to anonymously receive minimal error and health information from your device, the Smart Call Home service 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.

You can have both services configured on your system at the same time, although configuring the Smart Call Home service provides the same functionality as Anonymous Reporting, plus customized services.

## Configure Anonymous Reporting

To configure Anonymous Reporting, perform the following steps:

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Smart Call Home**.
  - Step 2** Check the **Enable Anonymous Reporting** check box.
  - Step 3** Click **Test Connection** to ensure that your system is able to send messages.  
ASDM returns a success or error message to notify you of test results.
  - Step 4** Click **Apply** to save the configuration and enable Anonymous Reporting.
- 

## Configure Smart Call Home

To configure the Smart Call Home service, system setup, and alert subscription profiles, perform the following steps.

### Procedure

- 
- Step 1** Choose **Configuration > Device Management > Smart Call Home**.
  - Step 2** Check the **Enable Registered Smart Call Home** check box to enable Smart Call Home and register your ASA with Cisco TAC.
  - Step 3** Double-click **Advanced System Setup**. This area consists of three panes. Each pane can be expanded or collapsed by double-clicking the title row.
    - a.** You can set up mail servers in the **Mail Servers** pane, through which Smart Call Home messages are delivered to e-mail subscribers.
    - b.** You can enter the information of the person to contact in the **Contact Information** pane for the ASA that appears in Smart Call Home messages. This pane includes the following information:
      - The name of the contact person.
      - The contact phone number.
      - The postal address of the contact person.
      - The e-mail address of the contact.

- The “from” e-mail address in Smart Call Home e-mail.
  - The “reply-to” e-mail address in Smart Call Home e-mail.
  - The customer ID.
  - The site ID.
  - The contract ID.
- c. You can adjust alert control parameters in the **Alert Control** pane. This pane includes the **Alert Group Status** pane, which lists the status (enabled or disabled) of the following alert groups:
- The diagnostics alert group.
  - The configuration alert group.
  - The environmental alert group.
  - The inventory alert group.
  - The snapshot alert group.
  - The syslog alert group.
  - The telemetry alert group.
  - The threat alert group.
  - The maximum number of Smart Call Home messages processed per minute.
  - The “from” e-mail address in Smart Call Home e-mail.

**Step 4** Double-click **Alert Subscription Profiles**. Each named subscription profile identifies subscribers and alert groups of interest.

- a. Click **Add** or **Edit** to display the **Subscription Profile Editor**, in which you can create a new subscription profile or edit an existing subscription profile.
- b. Click **Delete** to remove the selected profile.
- c. Check the **Active** check box to send a Smart Call Home message of the selected subscription profile to subscribers.

**Step 5** Click **Add** or **Edit** to display the **Add or Edit Alert Subscription Profile** dialog box.

- a. The **Name** field is read-only and cannot be edited.
- b. Check the **Enable this subscription profile** check box to enable or disable this particular profile.
- c. Click either the **HTTP** or **Email** radio button in the **Alert Delivery Method** area.
- d. Enter the e-mail address or web address in the **Subscribers** field.
- e. The **Alert Dispatch** area lets the administrator specify which type of Smart Call Home information to send to subscribers and under what conditions. There are two types of alerts, time-based and event-based, chosen according to how the alert is triggered. The following alert groups are time-based: Configuration, Inventory, Snapshot, and Telemetry. The following alert groups are event-based: Diagnostic, Environmental, Syslog, and Threat.
- f. The **Message Parameters** area lets you adjust parameters that control messages sent to the subscriber, including the preferred message format and the maximum message size.

**Step 6** For time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Add or Edit Configuration Alert Dispatch Condition** dialog box.

- a. Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:

- For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
  - For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
  - For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
  - For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.
- b. Click the **Basic** or **Detailed** radio button to provide the desired level of information to subscribers.
- c. Click **OK** to save the configuration.
- Step 7** For diagnostic, environment, and threat event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create or Edit Diagnostic Alert Dispatch Condition** dialog box.
- Step 8** Specify the event severity that triggers dispatch of the alert to subscribers in the **Event Severity** drop-down list, and then click **OK**.
- Step 9** For inventory time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create or Edit Inventory Alert Dispatch Condition** dialog box.
- Step 10** Specify how often to dispatch alerts to subscribers in the **Alert Dispatch Frequency** drop-down list, and then click **OK**.
- Step 11** For snapshot time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create or Edit Snapshot Alert Dispatch Condition** dialog box.
- a. Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:
- For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
  - For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
  - For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
  - For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.
  - For an interval subscription, specify how often, in minutes, the formation is sent to the subscribers. This requirement is applicable to the snapshot alert group only.
- b. Click **OK** to save the configuration.
- Step 12** For syslog event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create or Edit Syslog Alert Dispatch Condition** dialog box.
- a. Check the **Specify the event severity which triggers the dispatch of alert to subscribers check box**, and choose the event severity from the drop-down list.
- b. Check the **Specify the message IDs of syslogs which trigger the dispatch of alert to subscribers check box**.
- c. Specify the syslog message IDs that trigger dispatch of the alert to subscribers according to the on-screen instructions.
- d. Click **OK** to save the configuration.



- Step 13** For telemetry event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Telemetry Alert Dispatch Condition** dialog box.
- a. Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:
    - For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
    - For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
    - For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
    - For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.
  - b. Click **OK** to save the configuration.
- Step 14** Click **Test** to determine if the configured alerts are operating correctly.

## Monitoring Anonymous Reporting and Smart Call Home

See the following screens for monitoring Anonymous Reporting and Smart Call Home services:

- **Tools > Command Line Interface**

This pane allows you to issue various non-interactive commands and view results.

## History for Anonymous Reporting and Smart Call Home

**Table 40-1** History for Anonymous Reporting and Smart Call Home

| Feature Name        | Platform Releases | Description                                                                                                                                                                                                                                                                 |
|---------------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Smart Call Home     | 8.2(2)            | The Smart Call Home service offers proactive diagnostics and real-time alerts on the ASA, and provides higher network availability and increased operational efficiency.<br><br>We introduced the following screen:<br>Configuration > Device Management > Smart Call Home. |
| Anonymous Reporting | 9.0(1)            | You can help to improve the ASA platform by enabling Anonymous Reporting, which allows Cisco to securely receive minimal error and health information from a device.<br><br>We modified the following screen: Configuration > Device Management > Smart Call Home.          |

**Table 40-1** History for Anonymous Reporting and Smart Call Home (continued)

| Feature Name    | Platform Releases | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Smart Call Home | 9.1(2)            | The <b>show local-host</b> command was changed to the <b>show local-host   include interface</b> 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"> <li>• When a unit joins the cluster</li> <li>• When a unit leaves the cluster</li> <li>• When a cluster unit becomes the cluster master</li> </ul> <p>Each message that is sent includes the following information:</p> <ul style="list-style-type: none"> <li>• The active cluster member count</li> <li>• The output of the <b>show cluster info</b> command and the <b>show cluster history</b> command on the cluster master</li> </ul> |



## **PART 9**

### **Reference**





# Addresses, Protocols, and Ports

This chapter provides a quick reference for IP addresses, protocols, and applications.

- [IPv4 Addresses and Subnet Masks, page 41-1](#)
- [IPv6 Addresses, page 41-5](#)
- [Protocols and Applications, page 41-10](#)
- [TCP and UDP Ports, page 41-11](#)
- [Local Ports and Protocols, page 41-14](#)
- [ICMP Types, page 41-15](#)

## IPv4 Addresses and Subnet Masks

This section describes how to use IPv4 addresses in the Cisco 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.

## 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.

## Determine the Subnet Mask

See the following table to determine the subnet mask based on how many hosts you want.



### Note

The first and last number of a subnet are reserved, except for /32, which identifies a single host.

**Table 41-1** *Hosts, Bits, and Dotted-Decimal Masks*

| Hosts      | /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 |

## Determine 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.

### Class C-Size Network Address

For a network between 2 and 254 hosts, the fourth octet falls on a multiple of the number of host addresses, starting with 0. For example, The following table shows the 8-host subnets (/29) of 192.168.0.x.



#### Note

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.

**Table 41-2** *Class C-Size Network Address*

| Subnet with Mask /29 (255.255.255.248) | Address Range                |
|----------------------------------------|------------------------------|
| 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  |
| 192.168.0.16                           | 192.168.0.16 to 192.168.0.31 |

**Table 41-2**      **Class C-Size Network Address (continued)**

| Subnet with Mask /29 (255.255.255.248) | Address Range                  |
|----------------------------------------|--------------------------------|
| —                                      | —                              |
| 192.168.0.248                          | 192.168.0.248 to 192.168.0.255 |

## 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.
- The following table shows the 16 subnets of the network 10.1.



### Note

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.

**Table 41-3**      **Subnets of Network**

| Subnet with Mask /20 (255.255.240.0) | Address Range              |
|--------------------------------------|----------------------------|
| 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 |



# 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.

## Related Topics

[Configure IPv6 Addressing, page 15-12](#)

## 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: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). The following table shows several examples of address compression for different types of IPv6 address.

**Table 41-4** IPv6 Address Compression Examples

| Address Type | Standard Form               | Compressed Form        |
|--------------|-----------------------------|------------------------|
| Unicast      | 2001:0DB8:0:0:0:BA98:0:3210 | 2001:0DB8::BA98:0:3210 |
| Multicast    | FF01:0:0:0:0:0:0:101        | FF01::101              |
| Loopback     | 0:0:0:0:0:0:0:1             | ::1                    |
| Unspecified  | 0:0:0:0:0:0:0:0             | ::                     |



### Note

Two colons (::) can be used only once in an IPv6 address to represent successive fields of zeros.

An alternative form of the IPv6 format is often used when dealing with an environment that contains both IPv4 and IPv6 addresses. This alternative has the format `x:x:x:x:x: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.

## Unicast Addresses

This section describes IPv6 unicast addresses. Unicast addresses identify an interface on a network node.

### 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.

All global unicast addresses, other than those that start with binary 000, have a 64-bit interface ID in the Modified EUI-64 format.

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.

#### Related Topics

- [IPv6 Address Prefixes, page 41-10](#)
- [Interface Identifiers, page 41-8](#)
- [IPv4-Compatible IPv6 Addresses, page 41-7](#)

### 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-compatibility IPv6 address. The IPv6 transition mechanisms include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry a global IPv4 address in the low-order 32 bits. This type of address is termed an IPv4-compatible IPv6 address and has the format ::y.y.y.y, where y.y.y.y is an IPv4 unicast address.



### Note

The IPv4 address used in the IPv4-compatible IPv6 address must be a globally unique IPv4 unicast address.

The second type of IPv6 address, which holds an embedded IPv4 address, is called the IPv4-mapped IPv6 address. This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. This type of address has the format ::FFFF:y.y.y.y, where y.y.y.y is an IPv4 unicast address.

## Unspecified Address

The unspecified address, 0:0:0:0:0:0:0:0, indicates the absence of an IPv6 address. For example, a newly initialized node on an IPv6 network may use the unspecified address as the source address in its packets until it receives its IPv6 address.



### Note

The IPv6 unspecified address cannot be assigned to an interface. The unspecified IPv6 addresses must not be used as destination addresses in IPv6 packets or the IPv6 routing header.

## Loopback Address

The loopback address, 0:0:0:0:0:0:0:1, may be used by a node to send an IPv6 packet to itself. The loopback address in IPv6 functions the same as the loopback address in IPv4 (127.0.0.1).

**Note**

The IPv6 loopback address cannot be assigned to a physical interface. A packet that has the IPv6 loopback address as its source or destination address must remain within the node that created the packet. IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

## Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify the interfaces on a link. They need to be unique within a subnet prefix. In many cases, the interface identifier is derived from the interface link-layer address. The same interface identifier may be used on multiple interfaces of a single node, as long as those interfaces are attached to different subnets.

For all unicast addresses, except those that start with the binary 000, the interface identifier is required to be 64 bits long and to be constructed in the Modified EUI-64 format. The Modified EUI-64 format is created from the 48-bit MAC address by inverting the universal/local bit in the address and by inserting the hexadecimal number FFFE between the upper three bytes and lower three bytes of the of the MAC address.


For example, 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. The following figure shows the format of the IPv6 multicast address.

**Figure 41-1 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. The following table shows the prefixes for each IPv6 address type.

**Table 41-5**      *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

The following table lists the protocol literal values and port numbers; either can be entered in ASA commands.

**Table 41-6**      *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.         |

**Table 41-6 Protocol Literal Values (continued)**

| Literal | Value | Description                                                                                                               |
|---------|-------|---------------------------------------------------------------------------------------------------------------------------|
| 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.                                                                                          |

You can view protocol numbers online at the IANA website:

<http://www.iana.org/assignments/protocol-numbers>

## TCP and UDP Ports

The following table 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.

You can view port numbers online at the IANA website:

<http://www.iana.org/assignments/port-numbers>

**Table 41-7 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 |

**Table 41-7 Port Literal Values (continued)**

| <b>Literal</b> | <b>TCP or UDP?</b> | <b>Value</b> | <b>Description</b>                                                         |
|----------------|--------------------|--------------|----------------------------------------------------------------------------|
| 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 <b>exec</b> except that <b>cmd</b> 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 signaling                                                       |
| 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                                                                   |
| 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          | Mobile IP-Agent                                                            |



**Table 41-7 Port Literal Values (continued)**

| <b>Literal</b>    | <b>TCP or UDP?</b> | <b>Value</b> | <b>Description</b>                                                |
|-------------------|--------------------|--------------|-------------------------------------------------------------------|
| 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                                                    |
| 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

The following table 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 this table, 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 41-8** *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                       |
| 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)<br>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

The following table lists the ICMP type numbers and names that you can enter in ASA commands.

**Table 41-9** 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           |
| 31          | conversion-error     |
| 32          | mobile-redirect      |

